Mortgage valuation and the term structure of interest rates

MORTGAGE VALUATION AND
THE TERM STRUCTURE OF INTEREST RATES

MORTGAGE VALUATION AND THE TERM STRUCTURE OF INTEREST RATES

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de *Universiteit Maastricht*
op gezag van de Rector Magnificus, Prof. mr. G.P.M.F. Mols,
volgens het besluit van het College van Decanen, in het openbaar te verdedigen
op woensdag 8 december 2004 om 12.00 uur

door

Bart Henricus Matthijs Kuijpers

Promotores:
Prof. dr. ir. A.W.J. Kolen
Prof. dr. P.C. Schotman

Beoordelingscommissie:
Prof. dr. ir. C.P.M. van Hoesel (voorzitter)
Dr. W.F.M. Bams (Universiteit Maastricht)
Prof. dr. A.A.J. Pelsser (Erasmus Universiteit Rotterdam)

Mortgage Valuation and the Term Structure of Interest Rates

ISBN 90-9018761-8

Preface

The origin of this dissertation lies in the early spring of 1999, when Antoon Kolen offered me a Ph.D. position on mortgage valuation at the University of Maastricht. This position, being a joint project of the Department of Quantitative Economics and the Department of Finance, gave me the possibility to combine decision theory and finance. The combination interested me, and still does, and I started after finishing my study econometrics later that year. Although the last year of my Ph.D. has been difficult, I am still convinced that the choice I made five years ago was the right one at that time. Ultimately, it has resulted in this Ph.D. dissertation. I look back on a challenging and enjoyable period.

Several people contributed to this thesis. Many thanks go to my supervisors Antoon Kolen and Peter Schotman. Together, we analyzed many research problems from different perspectives. I am grateful to Antoon for sharing his knowledge and experience of operations research and his desire to solve mathematical problems. I thank Peter for passing on his financial and econometric knowledge and intuition. Also, I have very much appreciated the opportunity he gave me to work together in Stockholm on several parts of this dissertation, which led to a fruitful cooperation and a very enjoyable time.

Of the other places I had the privilege to visit during my Ph.D. period, Barcelona has made the most unforgettable impression. I am grateful to Bart, Carl and Roger for the great days in which we explored this beautiful city and the 'nearby' Pyrenees.

Of course I am thankful to my colleagues at the departments of Finance and Quantitative Economics for the pleasant atmosphere, inside and outside university. I have enjoyed our activities such as going to the movies and playing unihockey and soccer matches, even though I can hardly remember a won soccer game. Special thanks go to Frank for his inspiration, both during our study and the subsequent Ph.D. period.

Veruit mijn grootste dank gaat uit naar mijn ouders en broer, die er altijd zijn ls ik ze nodig heb. Hoewel jullie vaak niet precies wisten waar ik mee bezig was, toondenullie altijd belangstelling voor mijn vorderingen (en tegenslagen). Bij jullie vond, en vind inog steeds, de nodige afleiding. Zonder jullie steun zou dit proefschrift er niet zijn gewet.

Bart Kuijpers
September 2004

Contents

Chapter 1

Introduction

This thesis deals with mortgage valuation and interest rate tree calibration. Optimization and computation play a prominent role in both fields. Optimization is important for the derivation of a rational exercise policy of implicit mortgage prepayment options. Section 1.1 provides an overview of the Dutch mortgage market and describes typical Dutch mortgage features such as limited prepayment and tax issues. Since the term structure of interest rates is the main driver behind mortgage valuation and mortgage prepayment, interest rate modelling, interest rate tree calibration and interest rate derivative pricing are introduced in section 1.2. After discussing the modelling structures used throughout this dissertation, optimization aspects of mortgage valuation are introduced separately in section 1.3. Section 1.4 includes an outline of this thesis.

1.1 Mortgage valuation

A mortgage loan is a long term loan secured by a collateral, usually real estate. The mortgagor borrows money from the mortgagee and pays back the loan according to an agreed upon amortization schedule. In case the mortgagor fails to make the required payments, the mortgagee has the right to use the proceeds of the collateral to offset the loan, for example by selling the house.

The Dutch mortgage market has developed extremely fast from the early nineties on. An overview of the mortgage market in the Netherlands is provided by Alink [1], Charlier and Van Bussel [19] and Hayre [32]. Based on Charlier and Van Bussel and on data from CBS, the Dutch Central Bureau for Statistics, figure 1.1 shows that the total amount (in

FIGURE 1.1: Dutch mortgage market development. *Source: CBS*

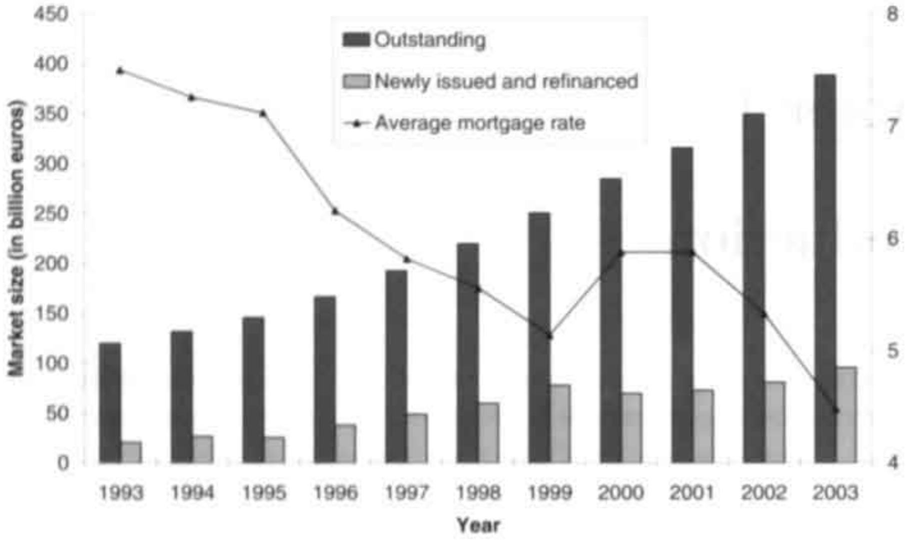

euros) of mortgage loans outstanding has more than tripled between 1993 and 2003. The proportion of newly issued and refinanced mortgages in the mortgage pool has increased, as the corresponding market share has more than quadrupled in the same period. This latter increase is mainly due to the rise of newly issued and refinanced mortgages in the years 1995 to 1999, a period in which the average mortgage rate dropped from 7.1% to 5.1%.

Annual mortgage transactions can be divided in issuing new mortgages and refinancing existing mortgage contracts. The amount of newly issued mortgages has hardly changed over the past ten years, according to figure 1.2. The increase in market share of newly issued and refinanced mortgage loans is completely due to refinancing existing loans, mainly driven by the significant mortgage rate decrease. Consequently, the importance of optimal interest rate driven prepayment and refinancing has increased. This dissertation covers both the derivation of optimal prepayment and refinancing strategies and the valuation of implicit prepayment and refinancing options.

Figure 1.3 shows the importance of mortgages on the combined balance sheet of Dutch banks. Mortgages make up for almost one quarter of the total bank's assets, which is more

FIGURE 1.2: Amount of newly issued and refinanced mortgages. *Source: CBS*

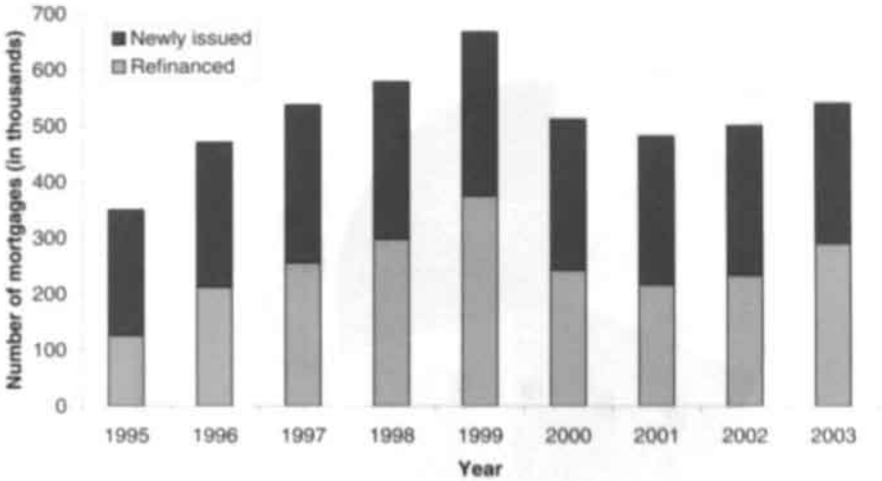

than bonds or short term loans and slightly less than long term loans.

Many types of mortgage contracts exist. A complete overview of Dutch mortgages in 2003 is available in the 'Hypothekengids 2003', the guide of the Dutch homeowners association 'Vereniging Eigen Huis'. A mortgage loan consists of several components. Loans may differ with respect to amortization schedule, contract rate adjustments and prepayment, refinancing or default options. Besides these basic ingredients of a mortgage contract, a variety of options is possibly included. Many contracts have insurance or investment opportunities. Also, tax regulations play an important role concerning the popularity of mortgage types.

Most commonly known amortization schedules include annuity mortgages and linear mortgages. Annuities are constant periodical payments including both redemption and interest. Initial payments are split into large interest payments and small redemption amounts. Later, when the remaining loan decreases, interest payments decline whereas redemption increases. Linear mortgages have constant amortization payments, but initially large total cash flows due to large interest payments.

A popular amortization schedule in the Netherlands is adopted by savings, investment or interest-only mortgages. With these mortgage types only interest payments occur during the lifetime of the contract. A savings mortgage is repaid at maturity, using a fund to

FIGURE 1.3: Total assets of Dutch banks. *Source: CBS*

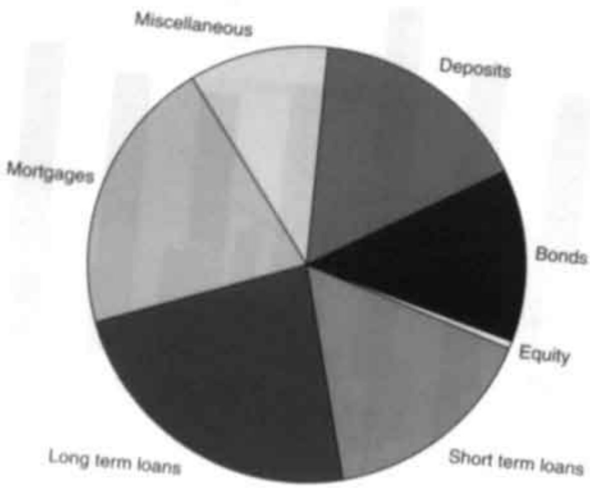

which payments have been contributed during the lifetime of the mortgage. The fund earns interest on a savings account. A more risky alternative is an investment mortgage, earning returns on bond and stock investments. An investment mortgage is exposed to uncertain future returns. Therefore the fund balance might not be sufficient to repay the full loan. An interest-only mortgage does not open a fund for repayment. Instead, the loan is repaid at maturity by private savings, taking out a new loan or selling the house.

The popularity of the above contracts is due to the Dutch tax regime. Since interest payments on mortgage loans are tax deductible, a high remaining loan (and thereby large interest payments and large tax deductions) is favorable. Because the income and tax situation is client specific, the advantage of this type of contracts is personal. Compared to other European countries, the Netherlands is the only country in which interest payments on a mortgage loan are fully tax deductible. Mainly for this reason the Netherlands is ranked second in Europe with respect to outstanding mortgage debt as percentage of GDP and outstanding mortgage debt per capita, according to Charlier and Van Bussel [19].

The development of the popularity of amortization schedules is confirmed by figure 1.4, based on Charlier and Van Bussel. In the early nineties annuity mortgages made up the

FIGURE 1.4: Market shares development of mortgage redemption schedules.

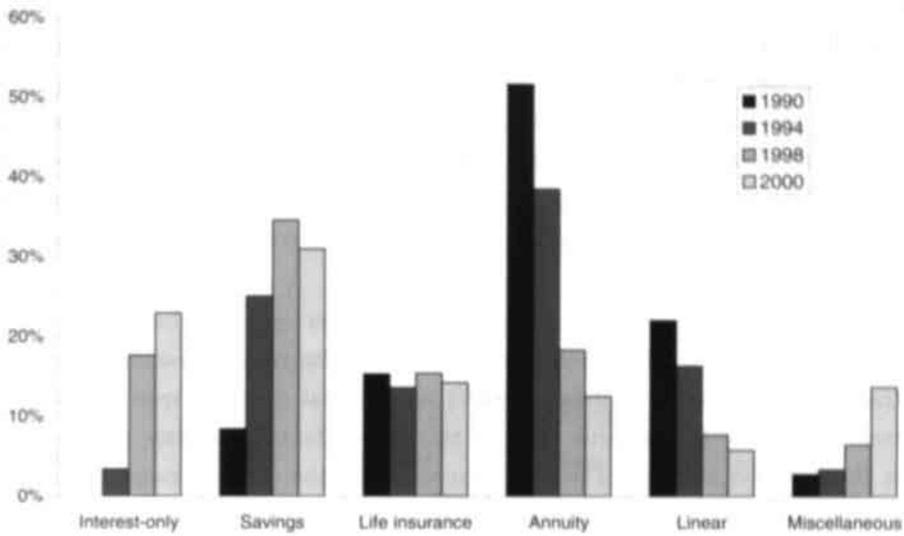

majority of the total mortgage pool. Nowadays the popularity of traditional redemption types (both annuity and linear mortgages) has decreased, in favor of interest-only mortgages and savings and investment mortgages, the latter type making up the largest part of the 'miscellaneous' category.

Mortgage contracts also differ with respect to fixed rate periods and contract rate adjustments. Longer fixed rate periods do not expose the borrower to future interest rate changes, but usually require a higher contract rate. A variable contract rate is attractive initially when the term structure is upward sloping, but implies a large borrower risk since any interest rate change is reflected in the contract rate. For longer maturity contracts, at the end of a fixed rate period the contract rate will be adjusted to match future interest rate conditions. With some contracts this adjustment is unrestricted. Others may have cap or floor restrictions to limit a contract rate increase or decrease respectively. In the Netherlands a typical mortgage contract has a fixed rate period of 5, 7 or 10 years, after which the contract rate is reset. The lifetime of a mortgage contract is usually 30 years. Thirty years is also the maximum period for tax deductions of interest payments.

A particular issue in the Netherlands is mortgage prepayment. While American mortgages can be fully called at any time, prepayment of Dutch mortgages is restricted. Each calendar year, prepayment of only a fixed percentage of the initial loan is allowed without penalty. This percentage depends on the type of contract and the bank at which the loan is taken out and usually equals 10, 15 or 20%. If the borrower decides to pay back the full loan at once, a penalty has to be paid which is equal to the sum of all present values of the future cash flow differences.[1] For some contracts even a threshold penalty exists, which might be larger than the prepayment gain. Due to the construction of the prepayment penalty no gain is possible from full prepayment, compared to prepayment of the maximally allowed percentage.

Besides rate adjustment or prepayment options, another option embedded in many mortgage contracts is a time for reconsideration. A time for reconsideration concerning the contract rate (a so-called 'rentebedenktijd') implies that during a specified interval of the fixed rate period (for instance the last two years) the contract rate can be fixed whenever the borrower chooses. The best moment to fix the contract rate is when the interest rate is lowest. When the borrower decides to exercise the reconsideration option, the contract rate equals the prevailing market rate for a new fixed rate period.

Low interest rates also give rise to prepayment and refinancing decisions. If the market conditions improved for borrowers who entered a mortgage contract when rates were relatively high, refinancing the contract or prepaying (part of) the loan in order to reduce interest payments might be favorable. However, except for the allowed prepayment percentage, transaction, administration or penalty costs involved can be larger than the expected gain. As a result, not every interest rate decrease will lead to refinancing or prepayment behavior.

Contrary to literature on American mortgages, default is of minor importance in the Netherlands. Every bank can check a national credit registry system before a mortgage is actually issued. Bad credits will face unfavorable borrowing conditions. Besides, the existence of a national mortgage guarantee (Nationale Hypotheekgarantie, NHG) decreases uncertainty for banks issuing new mortgage loans. For a mortgage contract including NHG, in case of borrower default, the guarantee foundation pays the remaining loan to the bank whenever the proceeds of selling the house are insufficient. The mortgagor is in debt to the foundation instead of the bank. Paying back the loan to the bank is assured. The

[1] Full prepayment of a mortgage is penalty-free when the mortgagor moves or dies and at a contract rate adjustment or refinancing date.

FIGURE 1.5: Mortgage rate vs. interest rates. *Source: CBS and DNB*

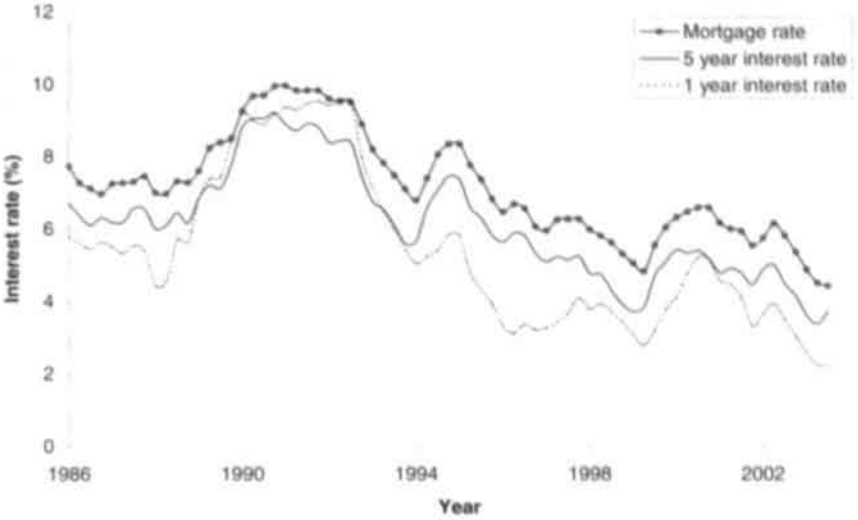

bank's risk therefore decreases and the contract rate will be lower compared to a mortgage contract without NHG.

We focus on the optimal prepayment strategy for mortgage loans, thereby dealing with mortgage valuation from a client's perspective. The role of a bank issuing mortgage contracts is to set the contract rate based on, among other aspects, prepayment behavior. The main contribution of this thesis to the existing mortgage valuation literature includes the valuation of the partial prepayment option and the derivation of the corresponding optimal prepayment strategy. We present various optimization algorithms, based on dynamic programming and linear programming, to obtain optimal mortgage values. An introduction to optimization issues will be provided in section 1.3.

For the valuation of mortgages, the development of interest rates is a key factor. The mortgage rate is highly correlated with (long term) interest rates, as can be concluded from figure 1.5, based on data from CBS and 'De Nederlandsche Bank' (DNB). Besides the interest rate level, interest rate volatility is important for the pricing of embedded options. Volatility is usually not observable, but implied by derivatives. Both the term structure of interest rates and volatilities implied by interest rate derivatives will be introduced in the next section.

1.2 Term structure of interest rates

The term structure of interest rates is the main driver behind mortgage prices and prepayment decisions. In mortgage valuation, prepayment can be accounted for in two ways. Either empirically observed prepayment or optimal, interest rate based prepayment is modelled. Empirically observed prepayments on Dutch mortgage contracts, for instance due to moving, have been analyzed by Alink [1] and Hayre [32].

Although borrowers might have different reasons to prepay a mortgage loan, our focus is on optimal, interest rate based prepayment. Optimal prepayment is interesting for both clients, minimizing the present value of all cash flows to amortize the mortgage loan, and for banks, to infer the impact on contract rates when mortgages are optimally prepaid. Optimal prepayment provides a worst case for mortgage issuers.

In this dissertation we will view mortgage loans as fixed income derivatives. No cash restrictions apply as soon as prepayment is optimal. If the cash position is insufficient to prepay (part of) the loan, we assume that the required amount can be borrowed at prevailing (lower) interest rates. Consequently, frictionless trading is assumed.

A term structure of interest rates represents current spot rates or prices of discount bonds (zero-coupon bonds) for varying maturities. Basic mortgage contracts, without prepayment or rate adjustment options, can be valued by the term structure only. Future cash flows are discounted using the spot rate with corresponding maturity in order to obtain the present value.

Prices of options on interest rate dependent assets, such as bond options, caps, floors, swap options (so called swaptions) and also mortgages including prepayment options, depend on interest rate volatility. To capture volatility the distribution of future interest rates must be known. Future interest rates are uncertain and can be modelled using different approaches and interest rate models.

In this thesis we will model uncertainty in future interest rates by using a discrete state space. A state space is a directed tree (of which a recombining and a non-recombining version are shown in figure 1.6). A node is also referred to as a state. The set of all states is partitioned into layers, each layer corresponds to one of the time points $t = 0, \dots, T$ and contains all the states that may occur at that time point. A state is represented by its layer t and index i as (i, t). Node $(0, 0)$ is called the root node. An arc connects two states in subsequent layers. Consider node k, indexed by (i, t), and node l, indexed by $(j, t + 1)$. If an arc (k, l) exists, then node k is called the predecessor of node l and l is called the

FIGURE 1.6: Discrete interest rate modelling: lattice versus tree.

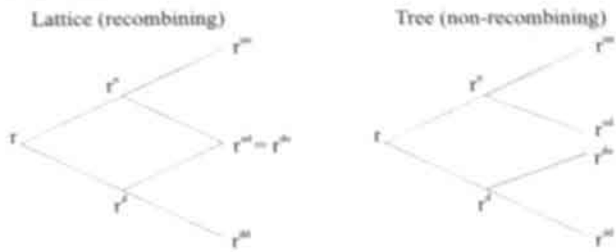

successor of k. Nodes without a successor are called leaf nodes. A node that is not the root node nor a leaf node is called an intermediate node. Associated with each arc (k, l) is a probability p_{kl} satisfying the property

$$\sum_{l:\exists(k,l)} p_{kl} = 1 \qquad \forall k. \tag{1.1}$$

So, p_{kl} is the probability that, given we are in state k at time t, we will be in state l at time $t+1$. An arc is referred to as a state transition. A path from the root to a leaf node represents one possible path of interest rates in the time interval $[0, T]$.

Scenario paths can be combined in various ways to form trees with different properties. Computationally efficient trees are not suitable for all types of problems. Some problems can only be solved using computational inefficient trees. In a binomial tree, for each non-leaf state at time t two possible states can be reached at time $t+1$. A non-recombining tree only contains scenario paths for which each state can be reached by exactly one path. An up movement followed by a down movement yields a different state (with a different interest rate or mortgage value) than a down movement followed by an up movement. Depending on the characteristics of the attribute[2] valued by the tree approach, in some special cases we may use a recombining tree, a so-called lattice. For computational purposes a lattice is much more efficient because the number of states at time t equals $t+1$, while the number of states in a non-recombining tree at time t equals 2^t. These two concepts are shown in figure 1.6. Interest rates can often be modelled by a lattice approach, but complex derivative pricing may require the use of non-recombining trees.

The valuation of derivative contracts relies on the absence of arbitrage opportunities.

[2]Possible attributes include shares, stock options and mortgages.

Since a mortgage is in essence a portfolio of elementary interest rate dependent contracts, taking out a mortgage is equivalent to investing in a (possibly complicated) bond portfolio. An investment strategy defines a portfolio for each non-leaf state consisting of zero-coupon bonds (that is, we sell and buy available zero-coupon bonds). To liquidate a portfolio we sell the assets that we own and buy back the assets we sold. For a non-recombining tree, the state contribution of an investment strategy is defined as the revenue of a portfolio if the state is the root state, is defined as the revenue of liquidating the portfolio of the unique predecessor state if the state is a leaf state (at the leaf node prices), and is defined as liquidating the portfolio of the unique predecessor state minus the cost of constructing the portfolio of the state itself if the state is an intermediate state.

An arbitrage opportunity is defined as an investment strategy for which every state contribution is non-negative and the sum of all state contributions is positive. If the contribution of the root node is positive, then we make a sure profit now without having any future costs. This is a sure way of making money. If the contribution of the root node is zero, then at least one price path exists for which the total contribution is positive. This situation is comparable with a free ticket in a lottery.

The existence of arbitrage can be formalized as follows. Let V denote the value process of an asset (or a portfolio of assets). Now $V(i,t)$ represents the value of the asset at time t in state i. An arbitrage opportunity, having root node contribution equal to 0, is then defined as a trading strategy such that

- $V(0,0) = 0$
- $V(i,t) \geq 0 \quad \forall i,t$
- $\exists (i,t) : V(i,t) > 0.$

Stated differently, an arbitrage opportunity is a possibility of making money, starting with nothing, without any risk of losing money.

A version of Farkas' Lemma shows that there are no arbitrage opportunities if and only if there exists a positive weight for each state transition such that the vector of prices at a state is the weighted sum over all successor states of the vector of prices at these states (for textbook references, see Duffie [27] and Pliska [73]). In a non-recombing tree, the state price of a state is defined as the product of all arc weights over all arcs on the unique path from the root to that state. By Farkas' lemma, the absence of arbitrage opportunities is equivalent to the existence of a vector of non-negative state prices. For a lattice, the state price of a particular state is the sum over all paths leading to that state of all state prices belonging to the same paths in the non-recombining tree.

One frequently used way to construct an interest rate model is to define in each node (i, t) the so-called short rate r_{it}, that is the interest rate for the time interval running from t to $t+1$. The positive arc weight of an arc rooted in (i, t) is defined as $\frac{1}{2(1+r_{it})}$. We assume equal up and down probabilities: $p_{ij} = \frac{1}{2}$ for both successor nodes. Dividing by $1 + r_{it}$, prices at time $t+1$ are discounted towards prices at time t. Once interest rates are defined for all states, a complete term structure can be derived in each node.

Given all arc weights it is easy to calculate all zero-coupon bond prices and show that the no-arbitrage conditions are satisfied. To see this, note that a zero-coupon bond has value 1 at its maturity. Using the arc weights we can calculate its value at the previous time point. Continuing this way, for each state the zero-coupon bond value at that state can be determined by multiplying the value at the successor node by the arc weight and summing the result over all successors.

A claim defines for each state a claim value. The interest rate model is complete if for every claim an investment strategy exists for which the contribution in each state is equal to the claim value. A necessary and sufficient condition for completeness is that for every non-leaf state the matrix with rows equal to the price vectors of all successor states has full row rank (see Duffie [27]). The interest rate model defined above is complete. One can view a claim as a financial product that pays the claim value if positive, and receives minus the claim if the claim value is negative. If the interest rate model is complete, then the price of the claim can be shown to be equal to the sum over all states of the product of the state price and the claim value.

A contingent claim is defined as a security having payoffs dependent (contingent) on the outcome of some underlying process (for instance, a price process of an underlying asset). As an example, consider an option paying out when expired in-the-money and not paying out when expired out-of-the-money. An Arrow-Debreu or state-contingent claim is a security paying 1 in one state at maturity, and zero in all other states. Denote the present value of a state-contingent claim paying 1 in state i at time t as $C_{00}(i, t)$. Many traded assets can be viewed as a portfolio of state-contingent claims. A bond maturing at time t pays 1 in every lattice state $i = 0, \ldots, t$ at t. Hence this bond is an equally weighted portfolio of the state-contingent claims having present value $C_{00}(i, t)$. The present bond price $P_{00}(t)$ follows as

$$P_{00}(t) = \sum_{i=0}^{t} C_{00}(i,t).$$

Interest rate trees are calibrated using an underlying term structure model. Thes models differ with respect to the number of factors and to the extent they capture futur drift, volatility and mean reversion of interest rates. For a given model an interest rat tree can be calibrated, such that model prices of interest rate dependent assets are as clos as possible to observed prices. Performance of a term structure model is measured as th difference between model prices and observed prices.

Both the future interest rate level and the future volatility or uncertainty are reflecte by observed market prices of interest rate dependent instruments. Bonds and swaps can b used to extract information about the level, whereas volatility is included in option data Implied volatilities are generally available from caps, floors or swaptions.

We will calibrate interest rates in order to match swap and swaption prices as closel as possible. A swap is a financial instrument to exchange a series of floating payments int fixed payments (or vice versa). The main use of a swap is to hedge financial risk present in future floating payments or revenues, for instance due to uncertain exchange rates in case of purchasing or selling goods in a foreign country. A swaption is the right, but not the obligation, to enter a swap contract at a certain date (the option expiration or exercise date) and a certain price (the strike or exercise price). Uncertain future interest rates determine the price development of swaps and swaptions. The current price must equal the sum of all discounted expected cash flows, both floating and fixed.

Optimal exercise of prepayment options in mortgage contracts is based on the volatility structure observed from swaptions. The next section introduces some of the optimization issues concerning prepayment decisions, related literature and an overview of optimization algorithms for the valuation of prepayment options applied in this thesis.

1.3 Optimal exercise of prepayment options

Exercise of prepayment options is based on the term structures of interest rates and interest rate volatilities. Much literature on Dutch mortgages, for instance Alink [1] and Charlier and Van Bussel [19], has focussed on empirical prepayment, which is not directly affected by interest rate driven prepayment decisions. In this dissertation, mortgage prepayment is triggered by the interest rate level. This approach provides a worst case for banks and an optimal prepayment strategy for clients.

Given a scenario tree of interest rates, optimal valuation of fully or partially callable mortgages can be modelled. The majority of literature on optimal exercise of prepayment options focusses on American mortgage contracts for which full and unrestricted prepayment is allowed without penalty. Optimal exercise policies for American mortgage loans have been derived by Kau, Keenan, Muller and Epperson [48]-[51]. A default option is typically included as well. Fully callable adjustable rate mortgages are also discussed. Hilliard, Kau and Slawson [37] apply a two-factor mortgage valuation model, the second factor being house price development.

Dutch mortgages, allowing only a limited prepayment amount, are more difficult to value, since partial prepayments imply path dependencies in the scenario tree. First, the remaining loan depends on earlier prepayments. Partially callable mortgages can have various remaining loan amounts. Second, the price of the (remaining) mortgage loan depends on the future prepayment strategy. Third, the calendar year restriction, constraining prepayment to a limited amount per year, imposes a restriction on the allowance of prepayments along parts of scenario paths belonging to the same calendar year.

Because of these path dependencies a mortgage value (and adopted prepayment strategy) can be different when reaching the same state, but having followed different paths. In general, valuation of partially callable mortgages requires the use of non-recombining trees. Since these are inefficient due to the exponential growth of the number of states, this thesis focusses on deriving lattice based algorithms to value partially callable mortgage loans.

Fortunately, some partially callable mortgages can be priced optimally by applying an efficient lattice approach, decomposing a mortgage contract into a portfolio of callable bonds. This approach is valid for mortgage loans that can be decomposed a priori, without knowledge of the optimal prepayment strategy. These mortgage contracts can be valued according to dynamic programming. Other mortgage types cannot be valued both optimally and efficient. For these contracts we derive a linear programming formulation.

1.4 Outline

The ordering of chapters in this dissertation describes the logical process, starting with observing interest rate derivative data, which are used for interest rate lattice calibration.

The resulting lattices, describing interest rate scenarios, are applied for mortgage valuation.

The first part of this thesis deals with the calibration of interest rate trees from observed data. For calibrating we require a term structure of interest rates, swaption prices and a term structure model. Chapter 2 provides an overview of several widely used term structure models. Characteristics, advantages and disadvantages of the models are discussed, while keeping in mind our purpose: to value long term mortgage contracts with typical embedded options.

Our data set includes swap rates, short term EURIBORs and implied swaption volatilities. Chapter 2 describes how to construct a term structure based on a spline method, for given swap rates and short interest rates. The valuation of both payer's and receiver's swaps is explained, given cash flow patterns and common quoting conventions. Black's option formula for (at-the-money) swaptions transforms swaption volatilities into swaption prices. Put-call parity shows that for at-the-money swaptions the prices for newly issued payer's swaptions and receiver's swaptions are equal.

In chapter 3 the data (term structure and swaption prices) and models of chapter 2 are used to calibrate a binomial interest rate lattice. A detailed technical analysis of the Black, Derman and Toy [9, BDT] model, as well as the *Ho and Lee [38, HL]* model, is provided. One of the most important characteristics of these models is the relation between volatility and mean reversion. These can not be matched independently, unless variable period lengths are allowed. The second part of chapter 3 is an extensive analysis of the calibration results. We consider input data on several dates and provide results on swaption pricing errors, term structure fitting, volatilities and mean reversion of interest rates.

Part II deals with mortgage valuation, based on the calibrated interest rate lattice resulting from part I. The typical Dutch prepayment feature of allowing a fixed percentage of the initial loan per calendar year introduces path dependencies in the binomial mortgage valuation tree. We focus on optimal prepayment behavior from a client's perspective.

Chapter 4 introduces distinctive features of mortgage contracts, including amortization schedules, call options and contract rate adjustments. We discuss mortgage valuation based on binomial lattice methods. Our focus is on deriving fair contract rates of common Dutch mortgage types. The fair rate is the contract rate for which the mortgage price is equal to the nominal loan value. For a mortgage quoted at the fair rate, neither bank nor client can make a profit. Fair rates are particularly useful when deriving option premiums as the difference between the fair rate of a mortgage including prepayment option and the fair rate of a similar mortgage excluding prepayment option.

Partial prepayment options may complicate valuation significantly because of path dependencies. Some mortgage types including partial prepayment options can be priced by using an efficient backward recursion algorithm, based on the number of calendar years and the number of prepayments remaining. This algorithm, dividing a mortgage loan into a portfolio of subsequently callable bonds, is developed in chapter 5 and can be applied to mortgage contracts for which the remaining loan amount only depends on the number of prepayments (interest-only, savings and investment mortgages, having no redemption payments during their lifetime), not on the periods in which these prepayments take place.

For general mortgage types, including both partial prepayments and a regular periodical amortization schedule, path dependencies cannot be removed in order to obtain the mortgage price efficiently. In chapter 6 we formulate a linear programming model for mortgages with partial prepayments. Optimality conditions are derived. Small instances can be solved based on a non-recombining tree, with an exponential number of states. For large instances a heuristic is required, providing an upper bound on the price and a lower bound on the fair rate. The dual formulation is used to obtain a lower bound on the price.

Chapter 7 compares fair rates of a large variety of mortgage contracts. Different amortization schedules and prepayment, rate adjustment and reconsideration options complicate the comparison of contracts. Which contract is cheapest, including all opportunities and restrictions, is not easy to determine. Fair rates are compared with observed mortgage rates and we evaluate the premiums of embedded options.

Mortgage prices and fair contract rates are determined for several variations of the underlying interest rate model to improve robustness. Results show that mortgage values are rather insensitive to the term structure model, the number of factors included and the length of the time steps in the pricing tree.

Finally, chapter 8 provides a summary of the results, concluding remarks and directions for future research.

Part I

Interest Rate Tree Calibration

Chapter 2

Term Structure Models and Data

2.1 Introduction

Modelling the term structure of interest rates is important for the valuation of interest rate dependent instruments such as bonds, swaps, bond and swap options, or mortgages. The purpose of this chapter is to introduce the ingredients from term structure models and data that are necessary for the calibration of interest rate trees in chapter 3 and, ultimately, for the valuation of mortgage loans in the second part of this dissertation.

The first important choice concerns the type of term structure model. Traditional models derive the term structure endogenously from assumptions on the dynamics of macro economic variables using equilibrium theory. These models derive a dynamic process for short term interest rates. Important aspects of the dynamics are drift, volatility and mean reversion. All other fixed income claims follow from no-arbitrage conditions. The best known of these models has been introduced by Cox, Ingersoll and Ross [23, CIR].

The main drawback of endogenous models is that they do not provide an exact fit for observed yield curves. As a result, the valuation of derivative securities is not accurate, since derivative prices are conditional on observed prices of plain bonds. To overcome this problem many term structure models have been extended. A dynamic process for the spot rate is constructed such that the implied yield curve is exactly equal to the observed yield curve. The extension involves time varying parameters that have to be re-calibrated every period. Since the observed yield curve is given, the extended models are called exogenous term structure models. We will review some of the well known models in section 2.2.

In order to model volatilities, one could construct a dynamic process for the spot rate that not only fits the observed yield curve, but also a set of liquid traded options. Other options will then be priced relative to the observed yield curve and the calibrated set of options. This will be the approach taken in this dissertation. We view a mortgage loan as a complex derivative security, which is priced relative to an observed yield curve and a set of observed option prices.

Since mortgage loans are modelled as fixed income securities, possibly involving complicated embedded options, valuation requires option pricing techniques. One of the powerful methods in no-arbitrage theory is risk neutral valuation. Under risk neutral valuation, expected future cash flows are discounted at the risk free short term interest rate. To justify the risk free rate as discount rate, the expectation is defined on a transformation of the original probability measure that governs the behavior of the spot rate. This new probability measure is called the risk neutral measure. In section 2.2 we will review the mechanics of the method. A detailed treatment and explanation is available in all major textbooks on option pricing.[1]

The second choice concerns the instruments on which the term structure model is calibrated. We use swap data to represent the term structure of interest rates, because options on swaps (so-called swaptions) are available to describe the volatility structure, the swap market is liquid for all maturities considered and the default risk of swaps is very limited (comparable to mortgages). Swaption data are used to model the term structure of interest rate volatilities. In sections 2.3 to 2.5 we discuss swaps and swaptions in detail and present the data. A method to transform raw swap data to a smooth yield curve of discount bonds is described. Observed swaption volatilities are transformed to swaption prices using Black's model. The yield curve and swaption prices obtained are used to evaluate the calibrated models.

Even with a preference for an exogenous term structure model and calibration to swaps and swaptions, there is still a wide range of candidate models. From a brief overview of recent empirical literature on swaption pricing in section 2.6 we conclude that a model that significantly outperforms all other models does not exist. All models have specific problems in fitting both the swap rate curve and a large set of swaptions. Combining different criteria (calibration, tractability, ease of implementation, possibility for generalization) we motivate our choice for a variation of the Black, Derman and Toy [9, BDT] model.

Although term structure models are presented in a continuous time setting, models are often discretized in applications with derivatives for which no closed form valuation

[1] See for example Hull [39], Duffie [27], Luenberger [59], Lyuu [60] and Rebonato [75].

formulas are known. The discretization we will use is a binomial tree or lattice, based on discrete time periods and discrete states. A tree is a set of scenario paths for which each state of the world can be reached by exactly one path. When pricing instruments in a discrete setting, optimality decisions (such as exercising an option) can be easily traced. For efficiency reasons we prefer to work with a lattice (a set of scenario paths for which different paths can lead to the same state), if possible. The basics of trees and lattices have been introduced in section 1.2. This chapter will be closed with a discussion of the implementation of the BDT model on a binomial lattice.

2.2 Overview of term structure models

Term structure models can be classified in many different ways. To start this overview we discuss several one-factor models. At the end of this section general frameworks will be considered. Hull [39], James and Webber [44], Pelsser [71] and Rebonato [75] give an extensive overview of interest rate models and their implications for calibrating term structures and pricing interest rate derivatives. This section provides a selective overview of the existing models and literature, a categorization of term structure models and model characteristics.

All term structure models are stated in continuous time as an Ito-equation, which takes a general form of

$$dr(t) = \mu(r,t)dt + \sigma(r,t)dz, \tag{2.1}$$

where $r(t)$ is the spot rate at time t, z is a standardized Wiener process with mean 0 and variance dt, and $\mu(r,t)$ and $\sigma(r,t)$ are the drift and the volatility measure of the spot rate, respectively. The majority of term structure models is defined in terms of the spot rate $r(t)$, although alternative formulations based on forward or swap rates exist as well.

Using the original probability measure we would need to risk adjust the discount rate depending on the risk of the cash flows. We apply the risk neutral probability measure in order to value uncertain cash flows by discounting the expected payoffs at the risk free rate, such that the risk neutral probability measure incorporates the risk adjustment. As a result, the drift parameter must be risk neutral (to discount future cash flows at the risk free rate). The adjustment of a general drift parameter $\mu(r,t)$ to a risk neutral drift parameter $\phi(r,t)$ will be discussed now. For more details one may read Hull [39], Ingersoll [43], Luenberger [59] or Lyuu [60].

Suppose that the zero-coupon bond price $P(r,t,T)$ follows

$$\frac{dP(r,t,T)}{P(r,t,T)} = \mu_P(r,t,T)dt + \sigma_P(r,t,T)dz, \tag{2.2}$$

where t is the issuing date of the bond and T is the maturity date. To obtain a risk free position we consider a short position in one bond maturing at T_1 and a long position in α bonds maturing at T_2. The parameter α will be chosen such that the bond portfolio is risk free at time t. The return on the portfolio equals

$$\begin{aligned} & -dP(r,t,T_1) + \alpha \cdot dP(r,t,T_2) \\ = & \ [-P(r,t,T_1) \cdot \mu_P(r,t,T_1) + \alpha \cdot P(r,t,T_2) \cdot \mu_P(r,t,T_2)]dt \\ + & \ [-P(r,t,T_1) \cdot \sigma_P(r,t,T_1) + \alpha \cdot P(r,t,T_2) \cdot \sigma_P(r,t,T_2)]dz. \end{aligned} \tag{2.3}$$

The bond portfolio stays risk free only if the weight α is updated continuously. For an instantaneously risk free portfolio, the volatility term must equal zero, hence

$$\alpha = \frac{P(r,t,T_1) \cdot \sigma_P(r,t,T_1)}{P(r,t,T_2) \cdot \sigma_P(r,t,T_2)}. \tag{2.4}$$

As an implication of no-arbitrage a risk free portfolio must earn the risk free rate r. Therefore the portfolio return must satisfy

$$\frac{-P(r,t,T_1) \cdot \mu_P(r,t,T_1) + \alpha \cdot P(r,t,T_2) \cdot \mu_P(r,t,T_2)}{P(r,t,T_1) + \alpha \cdot P(r,t,T_2)} = r, \tag{2.5}$$

stating that the absolute return (the drift term in equation 2.3) divided by the initial investment must be equal to the risk-free rate. Substituting for α and simplifying 2.5 leads to

$$\frac{\mu_P(r,t,T_1) - r}{\sigma_P(r,t,T_1)} = \frac{\mu_P(r,t,T_2) - r}{\sigma_P(r,t,T_2)} \equiv \lambda(r,t), \tag{2.6}$$

where λ, called the market price of risk, is independent of the bond maturity since T_1 and T_2 have been arbitrarily chosen. The instantaneous return on any asset depends on the asset's risk according to $r + \lambda(r,t) \cdot \sigma_P(r,t,T)$. A risk neutral process for the short rate is now represented as

$$dr(t) = \phi(r,t)dt + \sigma(r,t)dz, \tag{2.7}$$

where $\phi(r,t) = \mu(r,t) - \lambda(r,t) \cdot \sigma(r,t)$ is the risk free drift parameter. Using this drift allows us to discount future cash flows at the risk neutral probability measure. In the models discussed below, the effect of the short rate r on drifts and volatilities is included separately. Therefore we suppress the index r and write $\theta(t)$ for the model specific risk free drift and $\sigma(t)$ for volatility. Some models use a constant drift θ, a constant volatility σ, or both.

A selective overview of one-factor term structure models and their characteristics will be presented here. The main differences between the one-factor term structure models concern time dependency of the drift and volatility parameters, the degree of mean reversion and the impact of the interest rate level on the volatility term. We present the continuous time representation, although all models have an equivalent discrete version.

1. **Merton model**

   The Merton [65] model is specified by a constant drift parameter θ and a constant volatility parameter σ, yielding

   $$dr = \theta dt + \sigma dz. \tag{2.8}$$

   A significant drawback of this model is its inflexibility, due to the fact that both drift and volatility are independent of time. Also, $r(t)$ may become negative in some periods t. The Merton model implies negative long rates, because the short rate follows a random walk process with a constant drift and lacks mean reversion. Ingersoll [43] examines this effect in more detail.

2. **Ho and Lee model**

   The Ho and Lee [38, HL] model is the no-arbitrage version of the Merton model, allowing the drift parameter θ to be time dependent:

   $$dr = \theta(t)dt + \sigma dz. \tag{2.9}$$

   Still $r(t)$ can become negative. The drift parameter $\theta(t)$ is chosen to match the current term structure $P(t,T)$. Ho and Lee assume normally distributed short rates.

3. **Black, Derman and Toy model**

   Originally, the Black, Derman and Toy [9, BDT] model was introduced on a discrete state space. Subsequently, the continuous time limit has been derived. Following Hull [39], the model can be stated as follows:

   $$d\ln r = [\theta(t) + \frac{\sigma'(t)}{\sigma(t)}\ln r]dt + \sigma(t)dz. \tag{2.10}$$

The BDT model is a no-arbitrage model similar to HL. A significant advantage of BDT over HL is the model definition on the natural logarithm of the short rate instead of the short rate itself, preventing interest rates from becoming negative. This model definition implies that short rate volatilities are high when interest rates are high. Another strength of the BDT model is the inclusion of mean reversion. For decreasing volatility functions, interest rates in the BDT model exhibit mean reversion. The (logarithm of the) short rate decreases towards the long run average for large interest rates and increases for small rates. Black, Derman and Toy assume that short rates are lognormally distributed.

4. **Black and Karasinski model**

   The Black and Karasinski [10] model is similar to the BDT model, assuming a lognormal distribution of short rates, but allows for independent mean reversion and volatility:

   $$d\ln r = \kappa[\theta(t) - \ln r]dt + \sigma(t)dz. \tag{2.11}$$

   To capture mean reversion as well as drift and volatility, an additional degree of freedom is required, which is obtained by either using a trinomial lattice method or a binomial lattice with varying period lengths. To model three unknowns -drift, volatility and mean reversion- a binomial lattice (with constant period lengths) is not sufficient.

5. **Vasicek model**

   The Vasicek [79] model is an equilibrium model including mean reversion:

   $$dr = \kappa[\theta - r]dt + \sigma dz. \tag{2.12}$$

   Here $\kappa[\theta - r]$ represents the drift parameter. Both κ and θ are constant over time. The interest rate tends to move back to its natural average θ with rate κ. Volatility and mean reversion are modelled independently. Short rates are normally distributed.

6. **Hull and White model**

   The Hull and White [41] model can be seen as the exogenous version of the Vasicek model, including a time dependent drift parameter. The model is also similar to the Ho and Lee model, but including a mean reversion term:

   $$dr = \kappa[\theta(t) - r]dt + \sigma dz(t). \tag{2.13}$$

7. **Cox, Ingersoll, and Ross model**

   Some models include a positive correlation between the short rate and its volatility. The volatility of the short rate is large whenever the short rate itself is large. Also, the short rate volatility is small for low interest rates, implying that negative rates are unlikely. The Cox, Ingersoll, and Ross [23, CIR] model, an equilibrium model including mean reversion, developed in 1985, captures the positive correlation between interest rate level and volatility:

   $$dr = \kappa[\theta - r]dt + \sigma\sqrt{r}dz. \tag{2.14}$$

   CIR include the link between volatility and interest rate level explicitly, whereas the BDT model incorporates a similar effect due to the model definition on the natural logarithm of the short rate. Chan, Karolyi, Longstaff and Sanders [18, CKLS] concluded in 1992 that for a volatility term equal to σr^{γ}, $\gamma = 3/2$ provides the best fit.

No-arbitrage models are not exposed to arbitrage opportunities by construction, being set up from a martingale approach (using risk neutral probabilities, see Rebonato [75]). Models 1-7, defined on the short rate, contain the Markov property, stating that only the current state (and not the path to reach the state) affects the future conditional interest rate distribution. This justifies the use of interest rate lattices.

A general framework for many of the previously discussed term structure models (for instance Vasicek [79] and Ho and Lee [38]) has been introduced by Heath, Jarrow and Morton [34, HJM]. HJM allow for the inclusion of multiple factors, such that not all bonds of different maturities need to be perfectly correlated. Unlike the models discussed so far, HJM initially define a stochastic process for forward rates, instead of spot rates. The forward rate process is given by

$$df(t,T) = \mu(f,t)dt + \sum_{i=1}^{N} \sigma_i(T-t) \cdot f(t,T)dz_i, \tag{2.15}$$

where σ_i denote the volatility processes. The forward rate process in the HJM framework is non-Markovian. Current forward rates depend on the complete history of forward rates. Therefore, modelling forward rates requires a non-recombining tree, severely slowing down computations and limiting the number of periods that can be included.[2]

[2]For the HJM model, a lattice can only be used in case the volatility function belongs to a special class of volatility structures (see Li, Ritchken and Sankarasubramanian [55]).

For the definition of the original HJM forward rate process given by equation 2.15, forward rates are normally distributed, or equivalently, prices of zero-coupon bonds are lognormally distributed. This allows for negative forward and spot rates, and hence arbitrage opportunities when money can be stored without costs and risks. In case forward rates are assumed to be lognormal, negative rates are excluded. Unfortunately, interest rates might explode if these are continuously compounded, leading to zero prices for bonds and arbitrage opportunities.

Miltersen, Sandmann and Sondermann [66, MSS] introduced a framework in which simple interest rates over a fixed finite period are lognormally distributed. This framework, which became known as the LIBOR Market Model (LMM), was simultaneously developed by Brace, Gatarek and Musiela [13, BGM] and Jamshidian [45]. The lognormally distributed rates in LMM are consistent with the HJM framework for a specific choice of volatility, discussed by MSS. The forward rate process is given by

$$df(t,T) = \mu(f,t)f(t,T)dt + \sum_{i=1}^{N} \sigma_i(T-t) \cdot f(t,T)dz_i, \tag{2.16}$$

where σ_i denote the volatility processes and f faces simple compounding. For calibration purposes LMM has the same disadvantage as HJM: calibration to a binomial lattice is difficult because forward rates and swap rates are non-Markovian.

Until recently, LMM could be applied only for pricing European options, based on Monte Carlo simulation (see Rebonato [75]). At that time we did not consider LMM as a candidate term structure model for the valuation of mortgages with (American type) prepayment options. Recently, methods have been developed to suit LMM for pricing American options. For the first extensions of LMM, see for instance Andersen and Andreasen [3] and Longstaff and Schwartz [58]. Nowadays, market models are a serious alternative for pricing complex options. Many large investment banks currently use market models to value interest rate derivatives. Although we do not consider market models to price mortgages, the impact of a term structure model on mortgage valuation is analyzed in the second part of this thesis for robustness of the results.

The HJM framework and the LIBOR market model can naturally deal with multiple factors. Other multi-factor models have been introduced by Brennan and Schwartz [14] and Longstaff and Schwartz [57]. Brennan and Schwartz include a long term interest rate process as a second factor. They consider a stochastic process for the long rate and a process for the short rate oscillating around the long rate according to a mean reversion parameter. Longstaff and Schwartz [57] include a stochastic volatility process. Brigo and Mercurio [15] show that this model is equivalent to a two-factor extension of the CIR model.

The Black, Derman and Toy model can also be extended to a two-factor model, as will be done in chapter 3. Both factors are assumed to have all BDT properties, that is, they are lognormally distributed, face mean reversion, have non-negative interest rates and can be easily calibrated to a lattice. In the final sections of this chapter we will motivate our choice for the BDT model, based on model performance with respect to swap and swaption pricing. Before evaluating term structure models, we discuss the valuation of swaps and swaptions.

2.3 Notation for swap and swaption valuation

A swap is a financial instrument to exchange a floating leg and a fixed leg of payments, without exchanging the principal. The floating leg might be determined by floating interest rates, such as EURIBOR or LIBOR. The fixed rate determining the fixed leg of payments is called the swap rate. Note that in a swap contract usually only the net payments occur. Swaps are mainly used to hedge against uncertain payments or revenues in the future. The owner of a payer's swap pays a series of fixed amounts, while receiving floating cash flows. It can therefore be used when floating cash flows have to paid, to transfer these into a series of fixed payments, running less risk. Similarly, a receiver's swap can be used when facing positive floating cash flows, paying floating amounts in exchange for receiving fixed.

A swaption is an option to enter a swap at a certain time (the expiration date) and at a certain rate (the forward swap rate agreed upon when entering the swaption). Swaptions can be used for several purposes and are mainly an alternative to forward swaps. With a swaption, one might still profit from favorable interest rate movements, while being hedged against unfavorable movements. Contrary to forward swaps, swaptions will therefore have a positive price when settled, called the swaption premium.

In this section we introduce some notation for both swap and swaption valuation. The terms and concepts defined here will be explained in detail in the following sections. Concerning time issues we will adopt the following notation. The time unit is considered to be in years. A swap is entered at $t = \tau_0$, which can be either the current or any future period. A swap matures at its final period T. A swaption starts at t_0 and expires at t, when a swap can be entered by exercising the option. The conditions for this future swap to be entered are agreed upon at t_0. For this reason we will refer to t_0 as the agreement date of the swap. In case a swap is actually entered at t_0, then $t_0 = t$. A running index over time is usually represented by s. All time indices are annual.

During the lifetime of a swap, there are N payment dates: $\tau_i, i = 1, \ldots, N$. The last payment date equals the maturity of the swap, hence

$$t_0 \leq t = \tau_0 < \tau_1 < \tau_2 < \ldots < \tau_N = T.$$

The time s value of a swap entered into at t and maturing at T is denoted by $V(s), s = t_0, \ldots, T$, its swap rate agreed upon at t_0 by X. This swap rate is constant during the lifetime of the swap, although the future swap rate might change due to an evolving term structure of interest rates.

The principal of a bond with the same maturity as the swap is represented by B. The price of a zero-coupon bond with a lifetime from t to T equals $P(t, T)$. Zero-coupon bond prices will also be used to discount future cash flows and for defining the term structure of interest rates. Trivially, $P(t, t) = 1$. The term structure will also be represented as a yield curve, where the yield $y(t, T)$ is the $(T - t)$-period interest rate per annum. A third representation is a forward rate curve, with the one-period forward rates stated by $f(s, s+1) = r(s), s = t, \ldots, T - 1$[3]. Hence $r(s)$ is the forward rate over a period from s to $s + 1$. The relation between the three representations is given in section 2.4.4, where the resulting term structures are discussed.

The frequency or tenor of a swap is denoted by m and defined by the reciprocal of the number of cash flows per year. Typical tenors are 0.5 for semi-annual payments or 0.25 for quarterly payments. Payment frequencies for the floating leg and the fixed leg may be different. Common swap contracts in euros have semi-annual floating payments and annual fixed payments.

Because swaptions can be viewed as call or put options, their values are denoted by c and p. The implied volatility of the underlying swap is represented by σ, the underlying swap rate is again X, whereas the strike price is K. All swaptions considered are at-the-money when entered, therefore $K = X$ at initialization. The notation discussed here will appear frequently in the remainder of this chapter, together with less frequently occurring variables to be explained later.

[3]Note that in continuous time r denoted the instantaneous spot rate. In discrete time r denotes the one-period forward rate.

2.4 Term structure fitting using swaps

In this section we will construct term structures on selected days based on swap rates and short term EURIBORs. First, the availability of swap data is discussed.

2.4.1 Swap data

TABLE 2.1: EURIBORs and bid-ask averages of swap rates.

The table provides annualized EURIBOR data (in percentages) for February 29, 2000, February 15, 2001 and July 2, 2001. These short term interest rates have maturities for each month up to 1 year. Also, swap rates (in percentages) are included for the same dates as the average between bid and ask rate. Swap maturities range from 1 to 10 years.

EURIBOR	Feb 29, 2000	Feb 15, 2001	Jul 2, 2001
1 month	3.458	4.800	4.517
2 months	3.546	4.769	4.467
3 months	3.634	4.747	4.435
4 months	3.684	4.714	4.395
5 months	3.750	4.692	4.374
6 months	3.823	4.666	4.361
7 months	3.873	4.649	4.342
8 months	3.933	4.631	4.331
9 months	3.999	4.619	4.321
10 months	4.058	4.613	4.311
11 months	4.099	4.610	4.306
12 months	4.156	4.608	4.305
Swap rates	**Feb 29, 2000**	**Feb 15, 2001**	**Jul 2, 2001**
1 year	4.235	4.715	4.355
2 years	4.680	4.735	4.445
3 years	4.990	4.825	4.615
4 years	5.200	4.915	4.765
5 years	5.380	5.005	4.925
6 years	5.540	5.095	5.075
7 years	5.680	5.185	5.215
8 years	5.790	5.255	5.345
9 years	5.870	5.315	5.435
10 years	5.930	5.365	5.515

An example of a series of swap data at three arbitrary days (here February 29, 2000, February 15, 2001 and July 2, 2001) is provided in table 2.1. All listed swap rates are quoted against EURIBOR on a 30/360[4] basis. Floating payments occur twice a year, except for the 1 year swap which has a frequency of four payments per year. Fixed payments are annual. The swap rate is determined such that the initial value of a swap is zero. For example, a fair exchange between a floating leg and a fixed leg of payments occurs if a swap contract is settled to exchange EURIBOR to a fixed rate of 4.235% during one year, starting at February 29, 2000. At this swap rate, both parties in the swap agreement expect to break even.

In order to match the term structure in the first year, we use monthly EURIBORs, which are quoted on an act/360 basis and with the convention of transforming yields to prices by

$$P(t,T) = \frac{1}{1 + y \cdot (T - t)}. \tag{2.17}$$

Fabozzi [30] states similar conventions for US interest rates. The period $T - t$ (in years) is measured in actual days divided by 360. In case the payment day is a Saturday or a Sunday, the next Monday is considered to be the actual payment date, unless this Monday falls in the next month. In that case the Friday before is considered to be the actual payment date. The EURIBOR data for the three dates considered are provided in table 2.1.

For deriving a term structure of interest rates consisting of monthly periods the prices corresponding to swap rates and EURIBORs must be interpolated. To achieve this we will apply a spline method. A continuous, time-dependent function is fitted through the observed data, minimizing the sum of squared pricing errors. In the next subsection we first formalize the idea of the swap rate following from the fact that a swap contract is worthless when entered, thereby linking the swap rates to the term structure of interest rates. Then interpolation methods will be discussed to obtain a continuous zero-coupon bond price curve.

2.4.2 Pricing of swaps

For the valuation of swaps at the spot market, the swap starts as soon as the swap agreement is made (that is, $t = t_0$). As a first step in swap valuation we will price the cash

[4] An interest rate quoted on a 30/360 basis implies that each month is assumed to have 30 days and each year has 360 days. Other common quoting conventions include act/360 and act/act, where act indicates the actual days of a month or year.

FIGURE 2.1: The value and the payoffs of a floating rate bond.

The figure shows the payment dates of a floating-rate bond. Each payment date τ_i a cash flow of $r(\tau_{i-1})B$ is transferred. This amount equals the interest earned from $i-1$ to i on a bond with notational principal B. Consequently, the value of the bond is B just after each payment date.

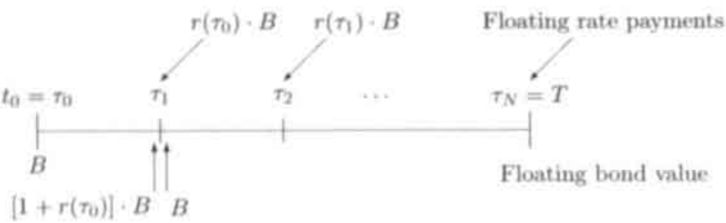

flows corresponding to the fixed and the floating leg separately. Fixed and floating legs can include the payment of a notational principal at maturity. This payment is equal for both legs and does not affect the swap value. Including principal, the fixed and floating legs are comparable to the cash flow patterns of a fixed-rate and floating-rate bond respectively. The floating-rate bond is worth B at initialization. At each payment date the interest earned in the previous period is paid, and the bond is again worth B immediately after each payment date (see figure 2.1). This can easily be derived from discounting the cash flows of an N-period floating-rate bond:

$$\begin{aligned}
B_{fl} &= \sum_{i=0}^{N-1} \frac{r(\tau_i)B}{\prod_{j=0}^{i}(1+r(\tau_j))} + \frac{B}{\prod_{j=0}^{N-1}(1+r(\tau_j))} \\
&= \sum_{i=0}^{N-1} \frac{(1+r(\tau_i))B - B}{\prod_{j=0}^{i}(1+r(\tau_j))} + \frac{B}{\prod_{j=0}^{N-1}(1+r(\tau_j))} \\
&= B.
\end{aligned}$$

The fixed-rate bond is worth the present value of all payments (which can be seen as zero-coupon bonds) plus the principal payment at maturity:

$$B_{fix} = X \cdot m \cdot B \cdot \sum_{i=1}^{N} P(t,\tau_i) + B \cdot P(t,T). \tag{2.18}$$

where X is the annual swap rate agreed upon at t_0 for a swap starting at t and maturing at the last payment date $T = \tau_N$, m is the frequency (e.g. 0.5 for semi-annual payments)[5], and $\tau_i, i = 1, 2, \ldots, N$ are the payment dates. The first term represents the periodical payments, the second term is the final principal payment. The time t value of a swap with a lifetime from t to T, $V(t)$, is the difference between the fixed leg and the floating leg of payments:

$$V(t) = B_{fix} - B_{fl} = X \cdot m \cdot B \cdot \sum_{i=1}^{N} P(t, \tau_i) + B \cdot P(t, T) - B, \tag{2.19}$$

where the floating leg is worth B initially by definition. The principal amount B can be scaled, resulting in an equivalent scaling of the swap value. We consider a receiver's swap here, that is the buyer of the swap receives the fixed leg of payments and pays the floating leg. A payer's swap has opposite value. At creation the swap has no value, so the swap rate (the fixed rate of the swap) X is set at a level at which the swap is worthless. The swap rate follows when setting $V(t)$ to zero:

$$X = \frac{1 - P(t, T)}{m \cdot \sum_{i=1}^{N} P(t, \tau_i)}. \tag{2.20}$$

After initialization the swap value may vary depending on the development of the floating rate. If this interest rate is lower (higher) than accounted for in the current term structure, then the payer of the floating rate gains (loses) and a receiver's swap has a positive (negative) value. The receiver of the floating rate loses (gains) and a payer's swap will have negative (positive) value. When calculating the one year swap rate from the EURIBOR data listed in table 2.1, these will not exactly match the observed swap rate. Reason for this is the day count convention. Swap rates are quoted at a 30/360 basis, whereas EURIBORs are listed at an act/360 basis. In case EURIBOR data are used to derive swap rates, the latter will also be on an act/360 basis. To obtain swap rates on a 30/360 basis, EURIBORs at a 30/360 basis are required. EURIBORs based on a 30/360 quotation will not be exactly the same as the rates used, resulting in a slightly different swap rate. In the sequel we allow for this small difference and use the 30/360 convention for the resulting monthly term structure of zero-coupon bond prices.

[5]Analogous to EURIBOR data, the payment frequency for swap rates must be adapted in case the payment date is a weekend day.

Pricing forward swaps

Up to this point only swaps traded in the spot market have been priced. The agreement to exchange fixed and floating legs was made at date t and the actual exchange occurred from date t on. Now the valuation of forward swaps will be discussed. Suppose that at time $t_0 < t$ we make an agreement to exchange fixed and floating legs, but the actual exchange only starts at time t. The fixed rate of the contract is determined at t_0, given the current term structure. Hence the swap is worthless at t_0, but might have a value when entered at time t.

Consider the current period to be t_0, the agreement date of a forward swap. Actual payments only start at time t. Given swap rate data we might infer the current term structure, that is, $P(t_0, s)$ for all $s = t_0, \ldots, T$. The value at time t of a forward swap with lifetime (t, T) can be directly inferred from equation 2.19. To obtain the current (time t_0) swap value we simply discount the swap value from t to t_0:

$$V(t_0) = X \cdot m \cdot B \cdot \sum_{i=1}^{N} P(t_0, \tau_i) + B \cdot P(t_0, T) - B \cdot P(t_0, t). \tag{2.21}$$

Solving for the swap rate by setting the swap value to zero yields

$$X = \frac{P(t_0, t) - P(t_0, T)}{m \cdot \sum_{i=1}^{N} P(t_0, \tau_i)}. \tag{2.22}$$

2.4.3 Term structure derivation

Based on the cash flow pattern of swaps a continuous term structure can be derived by applying the Nelson-Siegel [67] function or a spline method (following McCulloch [63]). Required data include annual swap rates and monthly EURIBORs.

The Nelson-Siegel function provides spot interest rates for each time period by

$$y(t, T) = \beta_0(t) + [\beta_1(t) + \beta_2(t)] \cdot (1 - e^{-\frac{T-t}{\gamma(t)}}) / \frac{T-t}{\gamma(t)} - \beta_2(t) \cdot e^{-\frac{T-t}{\gamma(t)}} \tag{2.23}$$

The parameters $\beta_0(t)$, $\beta_1(t)$, $\beta_2(t)$ and $\gamma(t)$ can be estimated to match observed swap (or bond) data. Main disadvantages of the Nelson-Siegel function are its inflexibility for short term interest rates and its impossibility to cope with (partly) decreasing yield curves.

A spline is a more flexible interpolation method to derive a continuous zero-coupon bond price curve, where prices depend on a polynomial or exponential function of the bond maturity. Spline methods are widely used, see for example Bams [6], Mathis and Bierwag [61] and Bali and Karagozoglu [5]. In order to obtain a continuous price curve of zero-coupon bonds we use a spline function on these prices, such that the sum of squared errors between the resulting swap prices (from 2.20 and 2.19) together with the short term bond prices (following from 2.17) and the observed data are minimized. Such a spline function might be polynomial (e.g. cubic) or exponential.

Exponential spline methods have the advantage that out of sample observations do not diverge when time approaches infinity. However, we consider a finite 10 year horizon of monthly periods. Cubic splines are more flexible because more parameters are used. Using a cubic spline method a number of breakpoints is chosen, dividing the time to maturity into several intervals, such that on each interval a cubic function with different coefficients can be used. Moreover, a cubic spline method involves performing a linear regression. For these reasons the cubic spline method is chosen to match the swap prices.

The general form of a cubic spline function to derive a price curve of zero-coupon bonds $P(t, T)$, with $\tau = T - t$ is given by

$$P(t,T) = 1 + \alpha_1 \cdot \tau + \alpha_2 \cdot \tau^2 + \alpha_3 \cdot \tau^3 + \sum_{l=1}^{L} \alpha_{3+l} \cdot [\tau - \beta_l]_+^3 \tag{2.24}$$

where $\beta_l, l = 1, \ldots, L$ are the breakpoints and $[\cdot]_+ = \max[.,0]$. $P(t,t)$ trivially has unit value. The time index τ varies continuously from t to $t + 10$ years. The number of breakpoints L is determined by a rule of thumb,

$$L = \lfloor \sqrt{M} \rfloor, \tag{2.25}$$

where M is the cardinality of the data set, that is, the total number of swap rates and EURIBOR data. According to our data set we are allowed to include four breakpoints, but we have used only three as we could not find a significant improvement with an extra breakpoint and we want to avoid overfitting.

To obtain the final regression we substitute the cubic spline equation 2.24 in equations 2.19 and 2.17 to obtain an expression for swap values and short term zero-coupon bond prices. Since the swap values are zero, the following joint regression must be performed:

$$m_j \cdot X_j \cdot \sum_{i=1}^{N_j} P(t, \tau_{ij}) + P(t, T_j) - 1 = u_j, \quad j = 1, \ldots, J \tag{2.26}$$

$$P(t, \tau_k) - \frac{1}{1 + y_k \cdot (\tau_k - t)} = u_k, \quad k = 1, \ldots, K \tag{2.27}$$

where u_j and u_k are the error terms, J is the number of swap rates and K the number of EURIBORs. This implies performing a regression on the following set of equations:

$$\begin{aligned}
u_j &= \sum_{i=1}^{N} m_j \cdot X_j + \alpha_1 \cdot (\sum_{i=1}^{N} m_j \cdot X_j \cdot [\tau_{ij} - t] + [T_j - t]) \\
&+ \alpha_2 \cdot (\sum_{i=1}^{N} m_j \cdot X_j \cdot [\tau_{ij} - t]^2 + [T_j - t]^2) \\
&+ \alpha_3 \cdot (\sum_{i=1}^{N} m_j \cdot X_j \cdot [\tau_{ij} - t]^3 + [T_j - t]^3) \\
&+ \sum_{l=1}^{L} \alpha_{3+l} \cdot (\sum_{i=1}^{N} m_j \cdot X_j \cdot [\tau_{ij} - t - \beta_l]_+^3 + [T_j - t - \beta_l]_+^3) \qquad (2.28) \\
u_k &= \alpha_1 \cdot (\tau_k - t) + \alpha_2 \cdot (\tau_k - t)^2 + \alpha_3 \cdot (\tau_k - t)^3 \\
&+ \sum_{l=1}^{L} \alpha_{3+l} \cdot (\tau_k - t - \beta_l)_+^3 + 1 - \frac{1}{1 + y_k \cdot (\tau_k - t)} \qquad (2.29)
\end{aligned}$$

Breakpoints are inserted after 1, 3 and 5 years, that is, $\beta_1 = 1$, $\beta_2 = 3$ and $\beta_3 = 5$. The sum of squared deviations to be minimized equals

$$\text{ssq} = \sum_{j=1}^{J} u_j^2 + \sum_{k=1}^{K} u_k^2. \qquad (2.30)$$

Having determined the zero-coupon bond price curve out of yearly swap rates and monthly interest rates according to EURIBOR, we will show the term structure of interest rates in the next section, in terms of prices, yields and forward rates.

2.4.4 Results

The spline coefficients and the resulting sum of squared errors to the regression stated in equations 2.28 and 2.29 are provided in table 2.2. The resulting term structures of interest rates based on the swap and interest data of table 2.1 are depicted in figures 2.2 to 2.4. Results are provided for February 29, 2000, February 15, 2001 and July 2, 2001.[6] For each figure, the top left diagram displays the zero-coupon bond price curve $P(t, T)$. We used

[6] A fourth instance, June 1, 2001, is considered as well. Results are similar to those of July 2, 2001 and are therefore not included in this chapter. In the second part of this thesis, results for June 1, 2001 are used for mortgage valuation.

FIGURE 2.2: Price, yield, and forward curves, February 29, 2000.

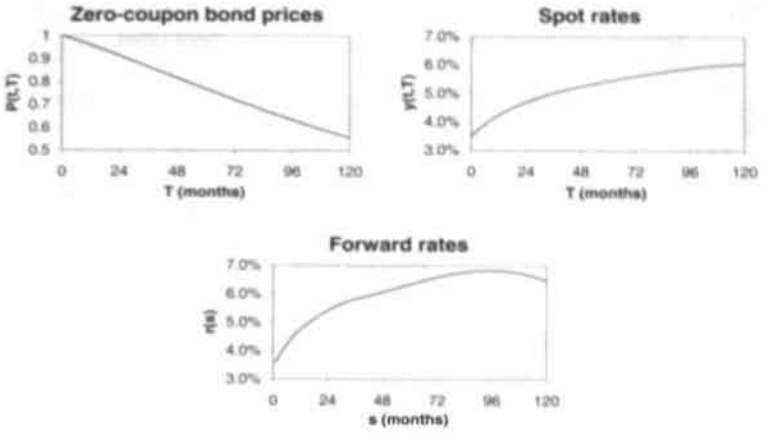

monthly periods spanning 10 years to fit the term structure to the swap rates and EURIBOR data.

A term structure can also be represented in terms of spot rates or yields. A spot rate $y(t,T)$ for a period starting at time t and ending at time T is defined as the annual $(T-t)$-period interest rate at time t. Therefore, the relation between zero-coupon bond prices and spot rates is given by

$$P(t,T) = \left[\frac{1}{1+y(t,T)}\right]^{(T-t)}, \tag{2.31}$$

or equivalently

$$y(t,T) = P(t,T)^{-\frac{1}{T-t}} - 1. \tag{2.32}$$

The yield curve $y(t,T)$ is depicted in the top right diagrams of figures 2.2 to 2.4 as a function of time T, consistent with the price curve $P(t,T)$ obtained before.

A third representation of the term structure is by using the forward curve containing all one-period forward rates over time. A forward rate (or short rate or one-period rate) starts at a future date s and ends at $s+1$. A spot rate $y(t,T)$ is then defined as the average of

FIGURE 2.3: Price, yield, and forward curves, February 15, 2001.

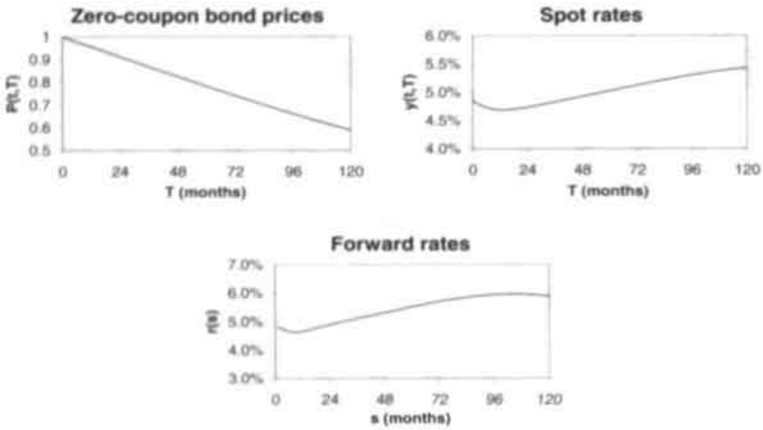

FIGURE 2.4: Price, yield, and forward curves, July 2, 2001.

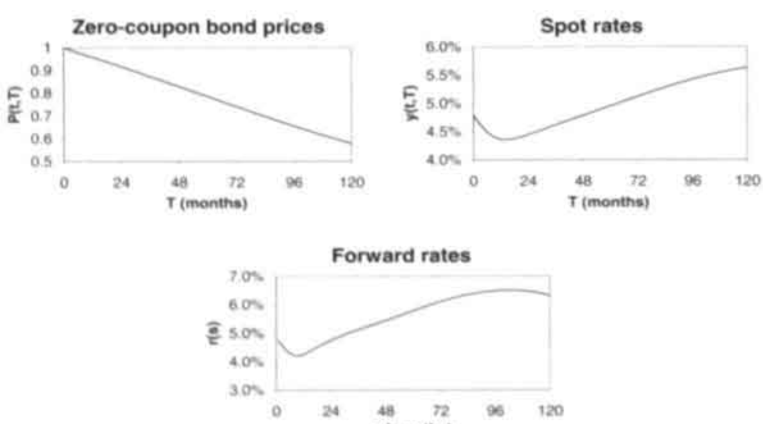

TABLE 2.2: Spline estimations and results.

The table provides the resulting regression coefficients of the joint regression stated in equations 2.28 and 2.29, for all three dates considered. The sum of squared pricing errors, including both swaps and short term zero-coupon bonds, is also reported. The principal of the swaps equals 1. Breakpoints are inserted after 1, 3 and 5 years.

Coefficients	Feb 29, 2000	Feb 15, 2001	Jul 2, 2001
α_1	-3.45 E-2	-4.74 E-2	-4.68 E-2
α_2	-7.97 E-3	3.96 E-3	8.35 E-3
α_3	1.85 E-3	-1.37 E-3	-3.33 E-3
α_4	-1.38 E-3	1.43 E-3	3.57 E-3
α_5	-4.79 E-4	-4.61 E-5	-2.18 E-4
α_6	1.18 E-4	4.40 E-5	1.04 E-4
ssq	3.75 E-7	2.92 E-7	3.82 E-7

the short rates $r(s)$ for all future periods $(s, s+1), s = t, \ldots, T-1$:

$$[1 + y(t,T)]^{T-t} = \prod_{s=t}^{T-1} [1 + r(s)]. \tag{2.33}$$

Short rates are derived from the spot rates by forward recursion. The resulting one-period forward rates are depicted in the bottom diagram of each figure.

The swap prices implied by the yield curve are presented in table 2.3 for February 29, 2000. Note that these swap values should be (close to) zero. To calculate the values, the observed swap rates of table 2.1 have been used. Besides, the swap rates are stated that imply a swap value of 0. The calculated swap rates differ between 0 and 0.6 basis points from the observed swap rates. A typical bid-ask spread for swap rates equals 3 basis points.

2.5 Swaption pricing

Besides the term structure to determine the interest rate drift, a volatility measure is required for interest rate tree construction. Interest rate volatilities are observed from options on interest rate swaps, so called swaptions. In this section we will outline the basics of swaption pricing using a Black and Scholes [11] approach.[7] In order to compare

[7]Longstaff, Santa-Clara and Schwartz [56] provide an accurate overview of the valuation of swaptions applying implied Black [8] volatilities.

TABLE 2.3: Implied swap values, February 29, 2000.

The table reports swap prices implied by the yield curve of February 29, 2000 and the swap valuation formula 2.19. Note that the swap values are zero when the term structure exactly matches the swap rates and EURIBOR data. The principal of the swaps equals 1. The swap rate that makes the swap value equal to zero is listed as well.

Maturity	Swap value	Swap rate	Maturity	Swap value	Swap rate
1	-1.50 E-5	4.237%	6	1.58 E-5	5.540%
2	1.03 E-4	4.674%	7	2.22 E-4	5.676%
3	1.57 E-4	4.984%	8	7.24 E-5	5.789%
4	-1.73 E-4	5.205%	9	-2.97 E-4	5.874%
5	-1.19 E-4	5.383%	10	1.21 E-4	5.928%

lattice prices of swaptions (obtained after calibration in chapter 3) we have to transform swaption volatility data into prices.

2.5.1 Swaption data

For calculating swaption prices two main ingredients are required: swap rates from the previous section and volatility data. Implied swap rate volatilities can be easily found in DATASTREAM. Swaption expiry dates range from one month to five years, after which a forward swap can be entered maturing between one and ten years. Swaption prices are quoted as implied Black volatilities. We used EURO vs. EURIBOR swaptions. Implied volatilities of selected swaptions are listed in table 2.4 for February 29, 2000. Figure 2.5 gives a graphical representation of all swaption volatilities at that date. Volatility is decreasing over time, but a 'volatility hump' is present both for small option maturities and for small swap maturities. This is a generally observed pattern, although the hump for other dates considered is less pronounced.

2.5.2 Black and Scholes method for swaption pricing

Swap rate volatilities are implied volatilities, being by no means the correct volatility parameters of interest rates, but only a different representation of the swaption price. Therefore we will calculate swaption prices both from the implied volatilities and from the interest rate tree. In chapter 3 will be explained how to match these prices. The price

TABLE 2.4: Implied swaption volatilities, February 29, 2000.

The table shows implied swaption volatilities (%) for a selection of swaption expirations and swap maturities on February 29, 2000. The data are retrieved from DATASTREAM. A typical bid-ask spread is 1 percentage point.

	swap maturities			
swaption maturity	1 year	2 years	5 years	10 years
1 month	16.05	17.00	15.70	14.40
3 months	16.35	16.70	15.35	13.60
6 months	17.05	16.45	14.75	13.00
1 year	17.70	16.90	13.90	12.25
5 years	15.90	13.80	10.35	9.30

of a swaption depends on the current and on the expected future yield curve, its option maturity, swap maturity, strike rate and the interest rate volatility.

Here we discuss the payoff structure of a swaption and show how to obtain swaption prices out of implied volatilities. Consider a swaption on a swap where we have the right to pay X_1 and receive a floating rate starting at the option maturity t and lasting $N = \frac{T-t}{m}$ periods of length m years. Such option on a payer's swap is called a payer's swaption. The principal is scaled to 1. Suppose that the (fair) swap rate at option maturity has changed to X_2. The payer's swaption is exercised if $X_2 > X_1$ resulting in a series of cash flows at $\tau_1, \ldots, \tau_N$, each equal to

$$m \cdot \max(X_2 - X_1, 0).$$

Hence the cash flows from this payer's swaption can be seen as payoffs from a call option on a forward swap with fair swap rate X_2 at option maturity and strike X_1 [8]. Analogously, the payoffs from a receiver's swaption equal

$$m \cdot \max(X_1 - X_2, 0).$$

Therefore a receiver's swaption can be viewed as a put option on the same swap with terminal value X_2 and strike X_1.

[8] Note that financial literature often refers to a payer's swaption as a put swaption. This can be explained by the fact that a payer's swaption is a put on a fixed-rate instrument. We follow Hull's notation [39] in the sequel, as payer's swaptions pay out for high future swap rates, and thereby resemble a call option on a swap rate.

FIGURE 2.5: Implied swaption volatilities, February 29, 2000.

The figure shows implied swaption volatilities of all available swaption expirations and swap maturities on February 29, 2000. The data are retrieved from DATASTREAM.

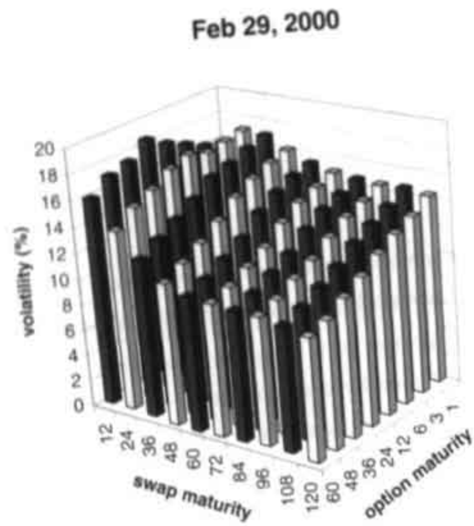

Now take a look at the payoff structure of a portfolio containing one payer's swap and one receiver's swaption (put option). The payoffs from the swap equal

$$m \cdot (X_2 - X_1),$$

resulting in portfolio payoffs equal to

$$m \cdot \max(X_2 - X_1, 0),$$

which is exactly the payoff structure from a payer's swaption (call option) with the same maturity and strike rate. When the swaption is entered, say at time $t_0 < t$, the swap rate X is determined at which the swap contract starts at time t. At time t_0 the swap is worthless, whereas it may have a value when the swap is actually entered at t. If we denote the initial value of the swap (at t_0) by $V(t_0)$, and the values of the call and put options on the same swap by c and p, we have $p + V(t_0) = c$. After all, if the payoffs from a payer's swap and a receiver's swaption (a put option) equal those of a payer's swaption (a call option), then both strategies must have equal costs to exclude arbitrage. Because $V(t_0) = 0$ (by definition at initialization), we have $c = p$ at the start of the option contract.

If the call and put prices are not equal we can make an arbitrage profit. Suppose $c > p$. By selling the call and buying the put and the payer's swap our immediate profit is $c - p > 0$. At option maturity the value of our portfolio equals zero, hence we make a total profit. Because the demand for the put increases, its price rises. Also, as the call is excessively supplied the price decreases, until $c = p$. Similar arguments hold for the case when $c < p$ to show that this situation cannot last either. As the call and put prices of an at-the-money swaption are equal, we may interpret the observed implied volatility as the price of both option types.

To conclude this section we present the Black-Scholes formula for the value c of a payer's swaption, giving the holder the right to receive a floating rate and to pay a fixed rate:

$$
\begin{aligned}
c &= m \cdot A \cdot [X \cdot N(d_1) - K \cdot N(d_2)] \\
&\quad \text{where} \\
A &= \sum_{i=1}^{N} P(t_0, \tau_i) = \sum_{i=1}^{N} [1 + y(t_0, \tau_i)]^{-(\tau_i - t_0)} \qquad (2.34) \\
d_1 &= \frac{\ln(\frac{X}{K}) + \frac{1}{2}\sigma^2 t}{\sigma\sqrt{t}} \\
d_2 &= \frac{\ln(\frac{X}{K}) - \frac{1}{2}\sigma^2 t}{\sigma\sqrt{t}} = d_1 - \sigma\sqrt{t}
\end{aligned}
$$

Here $y(t_0, \tau_i)$ is the annual interest yield for the period (t_0, τ_i) for discounting future swap payments to t_0 (time of entering the swaption), the swaption expiration or starting date of the swap is given by t, A is the total discounted payment during the lifetime of the swap, X is the underlying swap rate at t_0, K the exercise rate, m the tenor or frequency, and σ the underlying swap rate volatility. The principal is scaled to 1. Finally, $N(\cdot)$ is the standard normal distribution.

Formula 2.34 can be simplified a bit as all swaptions are issued at-the-money. When entering a swap *now* the swap rate is determined such that the initial swap value is equal to zero, according to equations 2.19 and 2.20. Also, when making a swap agreement now to enter a swap *in the future*, the initial swap value at the agreement date t_0 is zero and the swap rate is determined by equation 2.22.

Because both the swap rate and the exercise rate are determined such that the initial swap value is worthless, these rates are equal when making a swap agreement, whether payments start now or in the future. At swaption expiration however, the swap rate has changed due to a changing term structure over time and does not equal the exercise rate anymore. Hence, at option maturity the swap has some value on which the exercise decision is based. At the agreement date, by setting $K = X$ when entering a swaption and X given by 2.22, equation 2.34 simplifies to

$$\begin{aligned} c &= m \cdot A \cdot X \cdot [N(d_1) - N(d_2)] \\ &\text{where} \\ A &= \sum_{i=1}^{N} [1 + y(t_0, \tau_i)]^{-(\tau_i - t_0)} \\ d_1 &= \frac{1}{2}\sigma\sqrt{t} \\ d_2 &= -\frac{1}{2}\sigma\sqrt{t} = d_1 - \sigma\sqrt{t} \end{aligned} \tag{2.35}$$

The value of an at-the-money put (receiver's swaption) is

$$p = m \cdot A \cdot X \cdot [N(-d_2) - N(-d_1)], \tag{2.36}$$

with A, d_1 and d_2 as before. Trivially this results in $c = p$ for at-the-money swaptions.

Substituting for the swap rate given by equation 2.22, we observe that the market price of an at-the-money swaption does not depend on the tenor of the underlying swap:

$$c = [P(t_0, \tau_i) - P(t_0, T)] \cdot [N(d_1) - N(d_2)]. \tag{2.37}$$

Having transformed observed swap and swaption data into prices, the next section discusses recent literature on the performance of term structure models, in particular with respect to pricing swaps and swaptions. The swap and swaption prices obtained in the previous sections serve as a benchmark to evaluate model performance. A term structure model is accurate if model prices of swaps and swaptions are close to the benchmark.

2.6 Performance of term structure models

After observing or constructing a price curve of zero-coupon bonds, swaps can be exactly priced by the majority of term structure models. Volatility structures can usually not be modelled exactly. An extensive amount of literature exists on the performance of term structure models, interest rates being calibrated to both swaptions and caps and floors. We will focus on the calibration to swaption prices. Comments in Rebonato [75] are closely related to interest rate tree calibration on swaption prices observed at one particular date. Rebonato argues that prices of swaptions with the same maturity can be matched accurately using a one-factor model, but a one-factor model is considered unsuitable for pricing swaptions of various maturities because forward rate correlations implied by the swaption prices are not equal to 1 for different maturities.

Opposed to Rebonato, recent empirical results on historical estimation of swaption prices show that one-factor models perform quite well for several volatility structures, even compared to multi-factor models. Bühler, Uhrig-Homburg, Walter and Weber [16, BUWW] estimate the term structure on interest rate derivatives, especially warrants. They analyze the performance of seven one- and two-factor models, classified as either spot rate or forward rate models. BUWW conclude that a one-factor forward rate model with linear proportional volatility outperforms the other models, in terms of predictability. As a result, no clear evidence is found for the inclusion of multiple factors.

Driessen, Klaassen and Melenberg [26, DKM] and Fan, Gupta and Ritchken [31, FGR] apply several term structure models to price and hedge caps, floors and swaptions. They empirically investigate the term structure models to describe the behavior of these financial instruments over a period of several years, and predict future (in sample and out of sample) prices.

DKM apply option-based and interest-based estimation to model volatility. Option-based estimation takes derivative data into account, interest-based estimation models (co)variances of interest rates. Results indicate that for swaptions option-based models outperform interest-based models. DKM report best possible prediction errors for swaption pricing to be 8.5%, obtained by a Hull and White [41] based one-factor model. Increasing the number of factors does not improve accuracy: a three-factor model does not perform significantly better than a one-factor model. The swaps accurately fit the term structure, but all models considered tend to overprice swaptions with a short swap maturity.

The paper by FGR is closely related to DKM, but the results for swaptions in terms of pricing errors are somewhat more promising: average pricing errors are between 2.36% and 3.03% of the option price. As in DKM there is no need to include many factors to capture interest rate dynamics and price swaptions, since one- and two-factor parametric models perform at least equally well. FGR find, contrary to Rebonato, a better calibration performance on the pricing of swaptions, compared to cap and floor pricing. Swaption pricing errors are relatively small, but contrary to DKM, FGR observe an underpricing of swaptions with a lower swap maturity.

Longstaff, Santa-Clara and Schwartz [56] solve for the correlation matrix of forward rates implied by market prices of traded swaptions. Results indicate an overpricing of long maturity options and an underpricing of short maturity options. Average pricing errors are smaller than the assumed bid-ask spread of 6% for swaptions. Equivalently, pricing errors quoted in terms of implied volatility are below the bid-ask spread of 1 percentage point.

Peterson, Stapleton and Subrahmanyam [72] consider a multi-factor extension of the Black and Karasinski model, including factors describing forward rate behavior. Volatilities and mean reversion are assumed to be constant. Swaptions with a long option lifetime are underpriced, options expiring in the near future are overpriced, contrary to Longstaff, Santa-Clara and Schwartz [56]. Notice the striking inconsistencies between the papers discussed here with respect to pricing error patterns.

Note that our calibration of interest rate lattices, to be discussed in chapter 3, does not depend on underlying correlation structures or on historical option prices. All current option prices are only affected by the current interest rate lattice, not by past interest rate developments. This largely differs from the majority of the previously discussed literature, where calculated option prices (and thereby model performance) are based on model parameters estimated on past option prices. Besides a volatility function, these articles include historical price information for the valuation of current options.

Considering recent literature, including multiple factors to calibrate the term structure of interest rates on swap, swaption, bond, bond option, cap and floor prices does not significantly improve model performance. More factors might also lead to overfitting on one type of financial instrument, while out-of-sample assets are priced worse. Especially the results of Fan, Gupta and Ritchken [31] imply that for the accurate pricing of swaptions, correlations of forward rates do not necessarily have to match historical correlations. Given this empirical evidence, the next chapter analyzes calibration based on one- and two-factor variations of the Ho and Lee [38, HL] model and the Black, Derman and Toy [9, BDT] model, defined on a discrete state space.

The term structure models by BDT and HL can be captured in the same framework, but we will analyze the BDT model in more detail. BDT include empirically observed interest rate characteristics such as lognormally distributed and non-negative short rates and mean reversion. Mean reversion is obtained when volatilities are decreasing, which is possible but not enforced for our volatility specification in chapter 3.[9]

From the model definition 2.10 follows that the short rate volatilities underlying the BDT model are independent of (the logarithm of) the interest rate level. The one period volatility of the logarithms of the interest rates in two successor states does not depend on $\ln r$ (which is a necessary condition for trees to recombine). Consequently, the volatility of the short rate itself is larger for high interest rates and smaller for low interest rates. This feature is also incorporated in the Cox, Ingersoll and Ross [23] model and empirically observed by Chan, Karolyi, Longstaff and Sanders [18]. The Ho and Lee model assumes volatilities independent of the interest rate level. This partly explains the underperformance of the HL model compared to the BDT model, as found by Mathis and Bierwag [61].

2.7 BDT and binomial trees

This section introduces some properties of the Black, Derman and Toy [9] model on a discrete state space, thereby providing a preface to chapter 3, where a detailed specification is discussed. In a binomial tree or lattice, BDT assume that interest rate volatility only depends on time, and hence is constant for all states of the world in each period. Consider a one-period tree. The future one-period interest rate is either r_u or r_d, both with probability one half. The average (logarithm of the) short rate equals

$$\mu = \frac{\ln r_u + \ln r_d}{2}$$

and the variance of $\ln r$ is

$$\sigma^2 = \frac{(\ln r_u - \mu)^2}{2} + \frac{(\ln r_d - \mu)^2}{2} = \frac{1}{4}\left(\ln \frac{r_u}{r_d}\right)^2.$$

[9] In section 2.2 we have evaluated LIBOR market models. Both the performance and applicability of these models are promising. Only recently, market models are extended in order to deal with complex American type options. See Hull [39] for an overview.

Hence the one-period volatility equals $0.5 \cdot \ln \frac{r_u}{r_d}$. For multi-period trees the short rate average and volatility can be derived similarly by considering the two possible multi-period yields after one period.

A BDT based interest rate lattice must satisfy the no-arbitrage property. Section 1.2 has introduced the concept of arbitrage as a possibility of gaining money, without risk of losing money. By construction of a binomial tree or lattice, arbitrage is not possible when using risk neutral probabilities for upward and downward movements (see Pliska [73]). Without loss of generality, these risk neutral probabilities are assumed to be one half for each transition or arc in the tree. Every period, three parameters (drift and volatility of the BDT model and the risk neutral probability) have to be estimated, matching only two observations (term structure and volatility). Hence we choose to fix the probabilities to one half for all periods and states of the world, and calibrate the model to the observed yield curve and the volatility measure implied by swaption prices. Note that for trees to be recombining one degree of freedom is given up, as volatility parameters cannot be state dependent. The BDT model, to be calibrated on a binomial lattice, incorporates a volatility parameter which is only time dependent.

Some term structure models do not have the possibility of equating all probabilities to one half without loss of generality, but require more complicated lattices. An important difference between the BDT model and other mean reverting models is that BDT is defined on a random walk process, whereas for example the Hull and White model assumes a stationary state variable. For this reason, the HW model requires a more complicated lattice. A HW lattice reaches a maximum width; therefore the dynamics at the boundaries of the lattice must be different from the dynamics in the center. At some point the grid does not extend and at the edge there are sure upward or downward movements.

The BDT model is easily calibrated to a lattice. In the next chapter we calibrate swap and swaption prices and construct a lattice of short rates based on a variation of BDT. Additionally, we propose a two-factor extension of the BDT model, for which a trinomial lattice of interest rates is built.

2.8 Concluding remarks

This chapter has provided an overview of interest rate models and the valuation of interest rate derivatives. Several term structure models have been discussed in section 2.2. To calibrate an interest rate lattice to observed market prices of swaps and swaptions, the Black, Derman and Toy model will be applied in the next chapter. Its main characteristics include lognormally distributed and positive interest rates, positively correlated volatilities and interest rates, and mean reversion, although linked to volatility.

Empirically, one-factor models perform well compared to multi-factor models. Including more factors does often not improve calibration results significantly, while it may lead to overfitting of the in-sample assets. For instance, when the model is fitted to match caps, swaptions can still be priced poorly.

Within the class of one-factor models, the BDT model provides good calibration results. Interest rates are often assumed to be lognormally distributed, to have larger volatility when interest rate levels are higher and to be mean reverting. These features are covered by the BDT model. Also, mean reversion is sufficiently small to have an expanding tree in the distant future. Volatility decreases very slowly to 0. This appears to be no problem for our lattices with a 10 year horizon. For distant future time periods, when the number of states is sufficiently large to closely approximate a continuous distribution, a stable lognormal interest rate distribution is attained. In the next chapter we will analyze one- and two-factor variations of the BDT model for the calibration of interest rate lattices.

To evaluate model performance a term structure and swaption prices are required. Observed yearly swap rates and short term EURIBORs are transformed into a price curve for zero-coupon bonds by a cubic spline method. Besides a price curve, the term structure of interest rates can also be represented as a yield curve or a forward curve.

Swaption prices have been determined by transforming observed implied volatilities of at-the-money swaptions, applying the Black and Scholes [11] option model. This model can be rewritten in a format that allows the valuation of (at-the-money) swaptions. Payer's swaptions (call options on the swap rate) and receiver's swaptions (put options on the swap rate) are shown to have equal value, as long as both are at-the-money. Volatility data and calibration issues therefore hold for both payer's and receiver's swaptions.

The resulting yield curve and swaption prices generated in this chapter are used to evaluate the accuracy of the applied interest rate model when turning to calibration in chapter 3. Term structures of the three dates considered closely match swap rates and EURIBORs, as can be concluded from tables 2.2 and 2.3.

Chapter 3

Interest Rate Lattice Calibration[1]

3.1 Introduction

Binomial trees are widely used for pricing American and Bermudan type derivatives. In case of fixed income securities the nodes are related to spot interest rates. For practical use the tree is calibrated to the current term structure and usually the prices of a number of highly liquid European options, such as caps or swaptions. When the tree is calibrated it can be used to price American options by backward induction.

In constructing the tree there is a tradeoff between complicated trees that are calibrated on many instruments and simple trees that allow efficient computation of American options. Simple models have only one or two factors and use a recombining tree (lattice). In this chapter we are concerned with calibration of a lattice for subsequent use in pricing long term contracts with American options.

The literature contains a large number of one factor models that all have specific strengths and weaknesses.[2] We concentrate on extensions of some of the more popular models: Ho and Lee [38, HL] and Black, Derman and Toy [9, BDT]. Both are one-factor models that have a straightforward binomial lattice representation and can easily be calibrated to an initial term structure. In addition volatility parameters can be specified for both models such that they also exactly fit a series of caps or a subset of swaptions.

[1]This chapter is based on joint work with P. Schotman.

[2]The literature is so large that individual contributions would make up an excessively long list. Instead we refer to a number of textbooks that provide an extensive discussion of the theoretical and empirical properties of the many different one- and two-factor term structure models: Hull [39], James and Webber [44], Jarrow [46] and Rebonato [75].

In this chapter we consider the calibration to swaption volatilities in more detail. We search for the lattice that best fits a matrix of swaption prices with a range of option expiration dates and swap maturities. In calibrating a binomial tree one has a huge number of degrees of freedom, since in principle every spot rate at every node of the lattice can be chosen freely. In practice, to make calibration manageable, some structure is required, as in the BDT and HL models. Within either the HL or BDT model, we specify a flexible volatility function to provide the best possible fit for the swaption prices. As is well known from the literature this is not enough to price all swaptions exactly. Neither the BDT nor the HL model will deliver a perfect fit.[3] Swaptions are indexed by both the maturity of the underlying swap contract and the expiration of the option. This two-dimensional array of volatilities can not in general be fitted by a volatility function with a single time index.

The problem is due to the lack of control within the model over the correlation and relative volatilities of swap rates of various maturities. The literature suggests two ways to improve the model fit for the whole matrix of swaptions at a given time.[4] One solution is to add more factors, also with a flexible volatility specification. The other, partial, solution is to specify the degree of mean reversion of spot rates independent of the volatility like in the Black and Karasinski [10] and Hull and White [42] models. In the BDT and HL models the mean reversion properties of the spot rate are directly related to the volatility specification. Within a lattice we can change the mean reversion by making the time steps in the model time-varying. Like the extension to a two-factor model this complicates the subsequent application of the model for the valuation of other instruments. In the sequel we pursue both possibilities.

A recent alternative parameterization of the lattice is proposed by Peterson, Stapleton and Subrahmanyam [72]. They specify a two-factor dynamic model for the logarithm of spot rates with constant mean reversion and volatility parameters. Mercurio and Moraleda [64] construct a trinomial tree with a hump shaped volatility function.

We emphasize that this chapter is not meant to provide a test of term structure models or to obtain the best possible model for swaptions. Other models, like the LIBOR market model, are specifically designed for swaps and swaptions, and might very well fit swaption prices better than the lattices we construct in this chapter. Not being a lattice method,

[3]Backus, Foresi and Zin [4] provide an insightful analytical analysis of the pricing errors of calibrated HL and BDT models when the true process is a Vasicek [79] model.

[4]See Rebonato [75, chapter 3] for a textbook treatment and references.

the LIBOR market model, when accurately fitted to swap and swaption prices, is not immediately applicable to price a different class of financial instruments, especially American derivatives. Despite the advances in Monte Carlo techniques, a lattice remains the preferred tool for the valuation of complex derivatives where early exercise is important. Still, in evaluating the lattice models we will compare the patterns in the pricing errors with other models in the literature. Recent empirical studies investigating swaption pricing errors are Longstaff, Santa-Clara and Schwartz [56] and De Jong, Driessen and Pelsser [24].

The remainder of this chapter is organized as follows. Section 3.2 describes the construction of an interest rate lattice for a model with a flexible volatility structure and the extensions to general mean reversion or a second factor. The calibration method is discussed in section 3.3. Section 3.4 introduces input data for testing the specifications. Results are presented in section 3.5. Section 3.6 concludes.

3.2 Lattice construction

This section first explains the general framework. The next subsection takes up parameterization issues.

3.2.1 Framework

Our setup is related to Schmidt [78], but formulated in discrete time. Starting point for the construction of the lattice is a binomial random walk process for the state variable $S(t_n)$, where t_n is the time at the n^{th} step of the process. The state variable evolves according to

$$\begin{aligned} S(t_0) &= 0 \\ S(t_n) &= S(t_{n-1}) + U_n, \qquad n > 0, \end{aligned} \tag{3.1}$$

where the shock U_n can take two values,

$$\Pr[U_n = 1] = \Pr[U_n = -1] = \frac{1}{2}, \tag{3.2}$$

corresponding to either an up or a down movement. The initial time is $t_0 = 0$. Time steps have length $h_n = t_n - t_{n-1} > 0$. Time t_n is reached in n steps of length h_k $(k = 1, \ldots, n)$, so that $t_n = \sum_k h_k$. Changes in the state, $S(t_n) - S(t_{n-1})$, are uncorrelated over time, have zero mean and unit variance. The marginal distribution of the discrete process $S(t_n)$

is binomial with mean zero and variance n. At step n the state process $S(t_n)$ can take $n+1$ integer values between $-n$ and n with intervals of 2. Holding t fixed and increasing the number of steps n to reach $t = t_n$, the central limit theorem implies that $n^{-1/2}S(t)$ converges to a standard normal distribution.

Associated with the state process is a binomial lattice. If the process is at state $S(t_n) = i$, we are at node (i, n) in the tree. From node (i, n) the tree branches out to either nodes $(i-1, n+1)$ or $(i+1, n+1)$. At step n the lattice has $n+1$ nodes, indexed by i. The nodes run from $i = -n$ to $i = n$ with increments of 2.[5] The probability to reach state (i, n) is given by

$$\Pr[i, n] = \left(\frac{1}{2}\right)^n \binom{n}{(n+i)/2}. \tag{3.3}$$

At each node of the lattice we define the spot rates r_{in}. These are the spot rates at time t_n to discount values at t_{n+1}, the next step in the tree. The spot rate r_{in} is expressed on an annualized basis with discrete discounting over an horizon of length h_{n+1}. This gives the spot rates maturity h_{n+1}. The value at node (i, n) of one unit at nodes $(i+1, n+1)$ and $(i-1, n+1)$ is defined as

$$d_{in} = \frac{1}{(1 + r_{in})^{h_{n+1}}}. \tag{3.4}$$

We assume that the risk neutral dynamic process for the spot rates is related to the state process through the functional form

$$G(r(t_n)) = a(t_n) + b(t_n)\,\frac{S(t_n)}{\sqrt{n}}, \tag{3.5}$$

where $G(r)$ is a monotone increasing function. The spot rates in the tree are then

$$r_{in} = G^{-1}\left(a(t_n) + \frac{b(t_n)}{\sqrt{n}} \cdot i\right), \qquad i = -n, \ldots, n. \tag{3.6}$$

The function $G(r)$ determines the shape of the distribution of the spot rates. Popular choices for $G(r)$ are the identity $G(r) = r$ or the logarithmic $G(r) = \ln r$. The first leads to a generalization of the Ho and Lee [38] model, the second is a variation of the Black, Derman and Toy [9] or Black and Karasinski [10] model.

[5] From here on we will suppress the qualifier 'with increments of 2' when there can be no confusion. This makes notation lighter than the use of $i = 2j - n$ $(j = 0, \ldots, n)$.

The constants $a(t_n)$ are used to calibrate the tree to the current yield curve. From equations 3.5 and 3.6 we immediately find that the distribution of spot rates at time t_n has a median equal to $G^{-1}(a(t_n))$. The function $a(t)$ therefore controls the *location* of spot rates (conditional on the current spot rate $r(0)$). If interest rates are mean reverting we expect this function to converge to a constant.

Since the normalized state process $S(t_n)/\sqrt{n}$ has unit variance at all steps in the tree, the function $b(t)$ determines the standard deviation or *scale* of the spot rate levels. For a mean reverting process we must have

$$\lim_{t\to\infty} b(t) = b. \tag{3.7}$$

On the other hand, if spot rates follow a random walk, we expect $b(t)$ proportional to $\sqrt{t}$ for large t. Information about $b(t)$ is obtained from option volatilities.

Taking first differences with respect to both sides of equation 3.5 gives

$$\begin{aligned} G(r(t_n)) &= G(r(t_{n-1})) + a(t_n) - a(t_{n-1}) + \frac{b(t_n)}{\sqrt{n}} S(t_n) - \frac{b(t_{n-1})}{\sqrt{n-1}} S(t_{n-1}) \\ &= a(t_n) + \sqrt{\frac{n-1}{n}}\,\frac{b(t_n)}{b(t_{n-1})} \left(G(r(t_{n-1})) - a(t_{n-1})\right) + \frac{b(t_n)}{\sqrt{n}} U_n. \end{aligned} \tag{3.8}$$

This is a first order autoregressive process with parameters that are deterministically time varying.[6] With appropriate choices for $b(t_n)$ and h_n we can control various aspects of the spot rate process.

To further analyze the dynamic properties of the process we normalize the drift by h_n and the shock by $\sqrt{h_n}$. Also, let $y = G(r) - a$, and rewrite the dynamic process as

$$y(t_n) - y(t_{n-1}) = -\kappa(t_n)\, y(t_{n-1})(t_n - t_{n-1}) + \sigma(t_n)\, \epsilon(t_n) \tag{3.9}$$

where

$$\begin{aligned} \kappa(t_n) &= \frac{1}{h_n}\left(1 - \sqrt{\frac{n-1}{n}}\,\frac{b(t_n)}{b(t_{n-1})}\right) \\ \sigma(t_n) &= \frac{b(t_n)}{\sqrt{n h_n}} \\ \epsilon(t_n) &= \sqrt{h_n} U_n \end{aligned} \tag{3.10}$$

[6] Deterministic means that the entire sequences $a(t_n)$, $b(t_n)$ and h_n are functions of data as of time $t_0 = 0$.

The *volatility* of the process is given by $\sigma(t_n)$. The volatility of $\epsilon(t_n)$ has been normalized to $\sqrt{h_n}$, which is similar to the volatility of a Brownian motion. *Mean reversion* is defined by the coefficient $\kappa(t_n)$. Mean reversion occurs if $\kappa(t_n) > 0$. For a random walk with equal time steps, having constant $\sigma(t_n)$, $\kappa(t_n) = 0$.

Both $\kappa(t_n)$ and $\sigma(t_n)$ are intimately related to $b(t_n)$. More insight in the relation between $\kappa(t)$ and $\sigma(t)$ is obtained in the special case that time steps are equal. With $h_n = h = \Delta t$, we have that $t_n = nh$ and the process simplifies to

$$\Delta y(t) = -\kappa(t)y(t-h)\Delta t + \sigma(t)\epsilon(t) \tag{3.11}$$

with

$$\begin{aligned} \sigma(t) &= b(t)/\sqrt{t} && (3.12)\\ \kappa(t) &= -\frac{1}{\sigma(t-h)}\,\frac{\sigma(t)-\sigma(t-h)}{h} && (3.13) \end{aligned}$$

The mean reversion $\kappa(t)$ is a function of the volatility. The continuous time limit ($h \to 0$) of the process is[7]

$$dy = \frac{\sigma'}{\sigma} y dt + \sigma dz, \tag{3.14}$$

where $\sigma' = \partial\sigma(t)/\partial t$. The derivative dz is defined as a random variable with mean zero and variance dt.

Mean reversion implies that $b(t)$ is a non-increasing (but positive) function of time for large t. Hence $b(t)$ approaches a constant, $\sigma(t) \to 0$ and the dynamics of the process become deterministic for large t. Despite all the flexibility of the tree, this is a theoretical drawback of this construction method. The property is well known in the literature, see for example Rebonato [75, chapter 12] in the context of the BDT model. It is an implication of attempting to convert the random walk state process $S(t_n)$ into a mean reverting process for interest rates.

Returning to the general specification, the length of the time steps h_n provides an additional degree of freedom by which we can gain independent control over mean reversion and volatility. The time interval h_n appears in the denominator of the volatility parameter $\sigma(t_n)$ in equation 3.9. By taking shorter time steps for large n we can have a convergence to a stationary distribution ($b(t) \to b$) and non-zero volatility at the same time.

[7] See for example Hull [39, chapter 23] for a textbook reference to this result.

The relation between h_n and $\kappa(t_n)$ becomes more explicit by rewriting $\kappa(t_n)$ in 3.10 directly in terms of $\sigma(t_n)$ and h_n,

$$\kappa(t_n) = \frac{1}{h_n}\left(1 - \frac{\sigma(t_n)}{\sigma(t_{n-1})}\sqrt{\frac{h_n}{h_{n-1}}}\right), \tag{3.15}$$

as in Black and Karasinski [10]. Note the difference with the original Black and Karasinski model, where mean reversion at t_n depends on the length of the subsequent time period, because the first difference in equation 3.8 is taken with respect to t_{n+1}. To obtain the desired mean reversion at step n of the tree, we need to choose a time step that satisfies equation 3.15. Models with varying time intervals involve a relatively large number of steps at distant future periods, making them less efficient (or requiring more heavy use of pruning low probability segments of the lattice to prevent highly inefficient oversampling of the tails). To value short term options, smaller time steps in the beginning of the tree would be preferred, but this does not satisfy equation 3.15.

Two-factor model

For the generalization to a two-factor model we first define a bivariate lattice analogous to the construction in Jarrow [46, chapter 12]. The two state variables $S_1(t)$ and $S_2(t)$ have initial conditions $S_1(t_0) = S_2(t_0) = 0$ and evolve as random walks,

$$\begin{aligned} S_1(t_n) &= S_1(t_{n-1}) + U_{1n} \\ S_2(t_n) &= S_2(t_{n-1}) + U_{2n} \end{aligned} \tag{3.16}$$

The innovations U_{1n} and U_{2n} have a joint distribution with three possible outcomes,

$$\begin{aligned} \Pr[U_{1n} = 1, U_{2n} = 0] &= \frac{1}{2} \\ \Pr[U_{1n} = -1, U_{2n} = -1] &= \frac{1}{4} \\ \Pr[U_{1n} = -1, U_{2n} = 1] &= \frac{1}{4} \end{aligned} \tag{3.17}$$

Mean, variance and covariance of the state variables follow as

$$\begin{aligned} E[S_1(t_n)] &= 0, \qquad & E[S_2(t_n)] &= 0, \\ V[S_1(t_n)] &= n, \qquad & V[S_2(t_n)] &= n/2, \\ E[S_1(t_n)S_2(t_n)] &= 0. \end{aligned} \tag{3.18}$$

The two state processes are uncorrelated. If $n \to \infty$ for constant $t_n = t$, the processes $n^{-1/2}S_1(t)$ and $n^{-1/2}\sqrt{2}S_2(t)$ both have a limiting $N(0,1)$ distribution. The advantage of this construction method is that only three nodes can be accessed from any given node. Among the various methods to construct a two factor lattice, this structure has a minimal expansion of the number of nodes.

Spot interest rates in the two-factor lattice are defined as

$$G(r(t_n)) = a(t_n) + b_1(t_n)\frac{S_1(t_n)}{\sqrt{n}} + b_2(t_n)\frac{S_2(t_n)}{\sqrt{n/2}}, \tag{3.19}$$

for deterministic functions $a(t)$, $b_1(t)$ and $b_2(t)$. Each period, three successor nodes are required to solve for these three unknown parameters. A trinomial lattice is sufficient to describe interest rate dynamics for both factors. An extension towards two binomial lattices (one for each factor) is not required as that would lead to four equations and still three unknowns. In general, this approach leads to an $(n+1)$-nomial tree for an n-factor model.

Nodes in the two-factor lattice are defined as triplets (i, j, n), where at step n the index i can take values between $-n$ and n (with increments of 2) and the index j takes values between $-(n-i)/2$ and $(n-i)/2$ (also with increments of 2). The total number of nodes at step n is $\frac{1}{2}(n+1)(n+2)$. The root of the lattice is $(0,0,0)$. The possible transitions from node (i,j,n) at step $n+1$ are

$$\begin{array}{cclcc} & \nearrow & (i+1, j, n+1) & \text{with probability} & \frac{1}{2} \\ (i,j,n) & \rightarrow & (i-1, j+1, n+1) & \text{with probability} & \frac{1}{4} \\ & \searrow & (i-1, j-1, n+1) & \text{with probability} & \frac{1}{4} \end{array} \tag{3.20}$$

The probability to reach state (i,j,n) is given by[8]

$$\Pr[i,j,n] = \left(\frac{1}{2}\right)^{(n+i)/2}\left(\frac{1}{4}\right)^{(n-i+2j)/4}\left(\frac{1}{4}\right)^{(n-i-2j)/4} \cdot \frac{n!}{((n+i)/2)! \cdot ((n-i+2j)/4)! \cdot ((n-i-2j)/4)!}$$

[8]The state probabilities hold for all reachable states (i,j,n). Not attainable states trivially have $\Pr[i,j,n] = 0$.

These state probabilities can be rewritten into a simpler format with $k = (n+i)/2$ and $\ell = (n-k+j)/2$:

$$\Pr[k,\ell,n] = \left(\frac{1}{2}\right)^n \binom{n-k}{\ell}\binom{n}{k}.$$

Interest rates at each node of the lattice at step n are given by

$$G(r_{ijn}) = a(t_n) + \frac{b_1(t_n)}{\sqrt{n}} \cdot i + \frac{b_2(t_n)}{\sqrt{n/2}} \cdot j. \tag{3.21}$$

The first three steps of the two-factor lattice are shown in figure 3.1. The tree is trinomial with nonstandard branching. Possible transitions are depicted in an (i,j)-surface in figure 3.2, where the node labels represent the time period n.

For ease of parameterization we define the scaled processes

$$\begin{aligned} X_1(t_n) &= b_1(t_n)S_1(t_n)/\sqrt{n} \\ X_2(t_n) &= b_2(t_n)S_2(t_n)/\sqrt{n/2} \end{aligned} \tag{3.22}$$

Dynamics of the two scaled processes are

$$X_i(t_n) - X_i(t_{n-1}) = -\kappa_i(t_n)X_i(t_{n-1})(t_n - t_{n-1}) + \sigma_i(t_n)\,\epsilon_i(t_n), \tag{3.23}$$

where $\sigma_i(t_n)$ and $\kappa_i(t_n)$ are similar to the expressions in 3.10 and the innovations $\epsilon_i(t_n)$ have variance equal to h_n, $i = 1, 2$.

3.2.2 Parameterization

The lattice will be calibrated to an observed yield curve. This will determine the function $a(t)$. Assuming we have a complete smooth curve of zero-discount yields, determining $a(t)$ conditional on all other parameters will be easy and we do not need to put any further structure on the function $a(t)$. A smooth discount curve is available from fitting a spline function for discount bond prices based on observed coupon bond prices or swap yields. For given volatility parameters $b(t)$, $a(t)$ is determined by matching the term structure exactly applying a forward recursion.

The other input data for the calibration is a matrix of swaption volatilities for various combinations of option expiration dates and swap maturities. Providing a perfect fit to this entire matrix will in general not be feasible within a one-factor model, no matter

FIGURE 3.1: Two-factor lattice as a trinomial tree

The figure shows the first three time steps of the bivariate random walk process $(S_1(t_n), S_2(t_n))$ in (3.16). Transitions with thick lines have probability $\frac{1}{2}$, thinner lines have probability $\frac{1}{4}$.

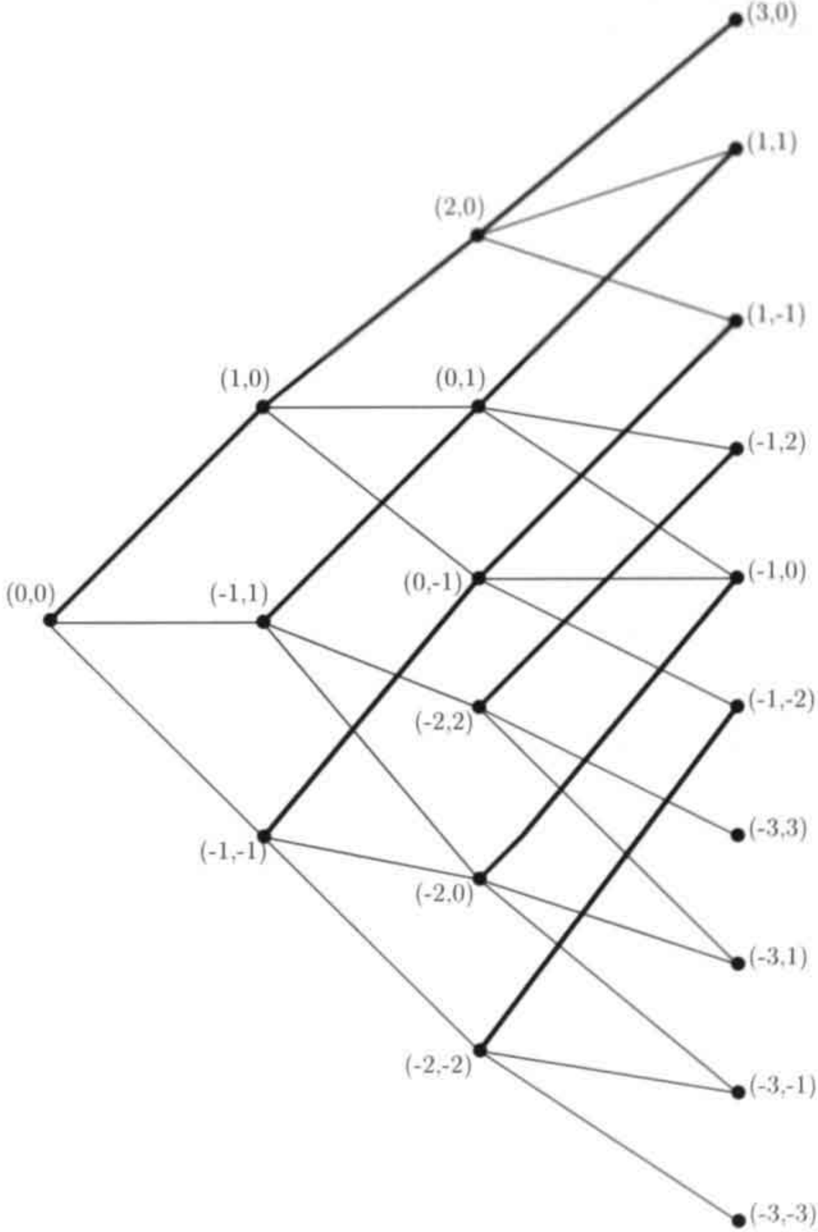

FIGURE 3.2: Two-factor lattice in (i,j) space

The figure shows the first four time steps of the bivariate random walk process (S_1, S_2) in (3.16). The process starts at the origin ($\odot$). The number below a node is the time step at which the node is reached. The process $S_1/\sqrt{n}$ evolves along the horizontal axis and $S_2/\sqrt{n/2}$ along the vertical axis.

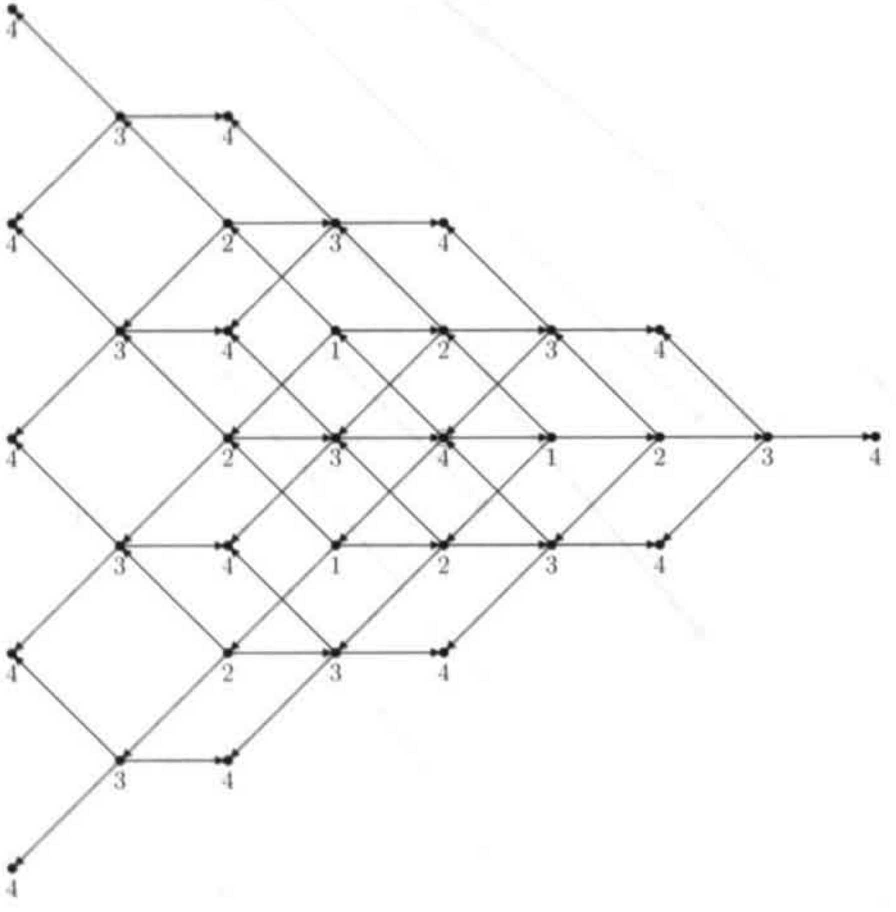

how flexible we are with respect to $b(t)$, h and $G(r)$. Similarly, leaving $b(t_n)$ and h_n as free parameters for all n leads to a huge number of parameters relative to the limited number of additional instruments. Such a large number of parameters will create numerical difficulties in finding the optimal fit and will also lead to considerable overfitting. For these reasons we need to impose conditions on the parameters.[9]

The analysis is essentially the same for any function $G(r)$. Since the input data are prices of European at-the-money options, the data will not be very informative about the distributional characteristics. The advantage of using $G(r) = \ln(r)$ is that all interest rates are positive by construction. The empirical advantage is that the actual empirical spot rate distribution is closer to a lognormal distribution than a normal. Another empirical advantage is that the volatility of r becomes proportional to the level of the spot rate. Chan, Karolyi, Longstaff and Sanders [18] and a large body of subsequent research provides ample evidence of the relation between volatility and the level of interest rates. As a robustness check all models are run with $G(r)$ either equal to r or $\ln r$.[10] In the remainder of the paper we will refer to models with $G(r) = r$ as Ho-Lee (HL) specifications and use BDT to denote the specification $G(r) = \ln r$.

The choice for h_n and $\sigma(t)$ are interrelated, since both jointly affect the volatility and the mean reversion of the process. Models with equal time steps $h_n = h$ are preferable from a computational point of view. Models with varying time intervals involve a relatively large number of steps towards the end of the tree. Time varying step sizes also need greater care in preparing the input data, since it must be possible to evaluate the yield curve at every maturity t. In the calibration stage every change in the parameters will lead to different points at which the yield curve is fitted. Time varying step lengths also do not guarantee that an option expires exactly at a node, a feature that is more important for subsequent uses of the lattice to price American options.

Unequal time steps are only required to correct for implausible mean reversion features of the calibrated process. Empirically, however, the restriction of equal period lengths does not create severe problems. We observe very little mean reversion in the actual time series behavior of short term interest rates. Time series models of interest rate dynamics usually

[9] Instead of restricting the model parameters, we could also interpolate the implied swaption volatilities in order to fill the entire surface $\sigma(\tau, T)$, with T being the swap maturity and τ the option expiration date. We have chosen to work with the fixed amount of hard data in order to reduce the numerical burden of having to search within a very high dimensional parameter space.

[10] Although the distributional assumptions may be less critical for fitting swaption volatilities, they are important for other instruments like out-of-the-money options.

do not provide precise estimates of mean reversion. Evidence for the risk neutral process comes from principal component analysis of yields, in which the first factor invariably has equal factor loadings at each maturity. Such a pattern of factor loadings is typically implied by a spot rate process that is close to a random walk.[11] Neither do we find much mean reversion implied by the relative volatility of long and short term interest rates. Both under the actual probability measure as under the risk neutral measure, mean reversion is likely to be positive but close to zero.

From equation 3.13 we know that mean reversion equals (minus) the rate of change of the volatility function. When the volatility $\sigma(t)$ is flat or slowly falling for large t, mean reversion $\kappa(t)$ will automatically remain close to zero. A small decrease in volatility is required to create enough mean reversion to make long term yields somewhat less volatile than short term interest rates. Under the same circumstances the implication that the spot rate process becomes deterministic for large t will not be a serious problem.

A special case with unequal time steps is the first order autoregression with constant volatility and constant mean reversion. With $G(r) = r$, this is the Hull and White [40] extension of the Vasicek model. For $G(r) = \ln r$, it is a restricted version of Black and Karasinski [10]. Setting $\sigma(t_n) = \sigma(t_{n-1})$ in equation 3.15, requiring $\kappa(t_n) = \kappa$, and solving for $\sqrt{h_n}$ gives

$$\sqrt{h_n} = \frac{1}{2\kappa}\left(\sqrt{\frac{1}{h_{n-1}} + 4\kappa} - \sqrt{\frac{1}{h_{n-1}}}\right) \tag{3.24}$$

as the implied time steps. The sequence h_n is strictly decreasing for $\kappa > 0$. Similarly, the limit as $\kappa \to 0$ leads back to the equally spaced lattice for a random walk with constant volatility. Decreasing step sizes as time evolves are unattractive, since the density of the lattice increases towards the end, whereas we would like to have finer step sizes early in the tree.

For models with time varying volatility there is an implicit dependence of $\sigma(t_n) = \sigma(t_{n-1} + h_n)$ on the step size. Using a first order approximation of $\sigma(t_n)$ in equation 3.15, we find the relation

$$\kappa(t_n) = \frac{1}{h_n}\left(1 - \frac{\sqrt{h_n}}{\sqrt{h_{n-1}}} + \kappa_0(t_{n-1})h_n\frac{\sqrt{h_n}}{\sqrt{h_{n-1}}}\right). \tag{3.25}$$

[11] Analytically this is seen most clearly in single affine models, where the limit as mean reversion goes to zero implies that yield curve movements will be parallel.

where $\kappa_0(t_{n-1}) = -\sigma'(t_{n-1})/\sigma(t_{n-1})$ is the implied mean reversion from the rate of change in the volatility. Rearranging gives

$$\frac{\sqrt{h_n}}{\sqrt{h_{n-1}}} = \frac{1 - h_n\kappa(t_n)}{1 - h_n\kappa_0(t_{n-1})}, \tag{3.26}$$

which has an approximate solution

$$h_n = \frac{h_{n-1}}{1 + 2(\kappa(t_n) - \kappa_0(t_{n-1}))}. \tag{3.27}$$

As long as the implied mean reversion $\kappa_0(t_{n-1})$ does not deviate too much from the desired mean reversion $\kappa(t)$ or from implausible large or small values, step sizes can remain constant.

Our main emphasis is, as in most of the literature, on the specification of the volatility function $\sigma(t)$. We conclude that the need for modelling $\kappa(t)$ is less than the need for modelling the volatility. We therefore explore in more detail models with a flexible functional form for $\sigma(t)$ with constant time steps h. Our benchmark model is the random walk with constant volatility and equal time steps $h_n = h$,

$$\sigma(t) = \sigma. \tag{3.28}$$

The resulting specification for $b(t)$ is

$$b(t) = \sigma\sqrt{t}. \tag{3.29}$$

We consider two more general types of functional forms for the volatility. The first is a function with exponential decay,

$$\sigma(t) = \left(\sum_{j=0}^{p} \theta_j t^j\right) \mathrm{e}^{-\kappa t}, \tag{3.30}$$

where p is the degree of a polynomial (equal to 1 in our applications). The implied functions for mean reversion and scale are

$$\kappa(t) = \kappa - \frac{\sum_{j=1}^{p} j\theta_j t^{j-1}}{\sum_{j=0}^{p} \theta_j t^j} \tag{3.31}$$

$$b(t) = \sqrt{t}\sigma(t) \tag{3.32}$$

The result for $\kappa(t)$ is the continuous time approximation $\kappa(t) = -\sigma'(t)/\sigma(t)$. The limits for $\kappa(t)$ as $t \to \infty$ or $t = 0$ are κ and $\kappa - \theta_1/\theta_0$ respectively. To avoid weird implications for the dynamics we assume that the polynomial $\sum_j \theta_j t^j$ does not have roots for positive t. Restricting $\theta_j > 0$ for all j is an easy way to impose this.

An important special case is $p = 1$,

$$\sigma(t) = (\theta_0 + \theta_1 t)\,\mathrm{e}^{-\kappa t}. \tag{3.33}$$

With equal time steps, volatility can be hump-shaped. For positive θ_0 and θ_1 the volatility reaches a peak at

$$\bar{t} = \frac{1}{\kappa} - \frac{\theta_0}{\theta_1}. \tag{3.34}$$

Due to the exponential decline of $\sigma(t)$, also the scale function $b(t)$ goes to zero. For large t the process converges to a single point.

The implausible long run implications of the exponential volatility model motivate the second functional form. We would like to allow for a rate of decline in the volatility function that is less than exponential. One functional specification that can have this property is

$$\sigma(t) = F(t)\frac{\beta_0}{\sqrt{t}} + (1 - F(t))(\beta_1 + \beta_2 t), \tag{3.35}$$

where the weighting function

$$F(t) = \frac{\alpha \cdot t^\theta}{1 + \alpha \cdot t^\theta}. \tag{3.36}$$

When $\theta > 0$ the weight function approaches 1 for large t and $F(0) = 0$. If $\theta > 1$ the volatility goes to zero for large t. For smaller values of θ the process exhibits explosive behavior. For values $1 \leq \theta < 3/2$, the volatility goes to zero, but the scale $b(t) = \sigma(t)\sqrt{t}$ still diverges. For the existence of a limiting variance of the spot rate we need the stronger condition $\theta \geq 3/2$. The convergence to the constant $b(\infty) = \beta_0$ is at the rate $t^{3/2-\theta}$, and thus extremely slow if θ is only slightly above the threshold. The initial volatility, $\sigma(0)$, is well defined if $\theta \geq \frac{1}{2}$. For $\theta > \frac{1}{2}$, volatility starts at β_1.

Two-factor model

Parameterization of the two-factor model is similar to the single factor model. The function $a(t)$ is used to fit the current yield curve, while the functions $\sigma_i(t)$ will be calibrated to implied swaption volatilities. The functional forms for both volatility functions are identical to equations 3.30 and 3.35 in the single factor case. Analogous to 3.35, the instantaneous volatilities of both factors are given by

$$\sigma_i(t) = F_i(t)\frac{\beta_{i0}}{\sqrt{t}} + (1 - F_i(t))(\beta_{i1} + \beta_{i2}t), \tag{3.37}$$

where the weighting functions $F_i(t)$ are

$$F_i(t) = \frac{\alpha_i \cdot t^{\theta_i}}{1 + \alpha_i \cdot t^{\theta_i}}. \tag{3.38}$$

Time steps in the two factor model are the same for both factors. As for the single factor model we will initially use constant time steps $h_n = h$, so that the mean reversion of both $X_i(t)$ processes is equal to the rate of change in the volatility.

3.3 Calibration

This section discusses the determination of the parameters $\sigma(t)$ and h_n. We start with the single factor model, than generalize to the two-factor model.

3.3.1 One-factor model

The short term interest rates at step n are

$$G(r_{in}) = a(t_n) + \frac{b(t_n)}{\sqrt{n}}i, \qquad i = -n, \ldots, n \tag{3.39}$$

with parameters $a(t_n)$ and $b(t_n)$, $n = 0, \ldots, N$, to be calibrated. For the calibration we start with a forward recursion to construct the Arrow-Debreu state prices q_{in} ($i = -n, \ldots, n$). The state prices q_{in} are the prices at $t = 0$ of elementary securities with payoff 1 at node (i, n) and zero at all other nodes. Start with $q_{00} = 1$. Given a tree up to step n, the state prices at step $n + 1$ in a binomial tree with transition probabilities $\frac{1}{2}$ can be derived as[12]

$$\begin{aligned} q_{-(n+1),n+1} &= \frac{1}{2}q_{-n,n}d_{-n,n} \\ q_{i,n+1} &= \frac{1}{2}\left(q_{i+1,n}d_{i+1,n} + q_{i-1,n}d_{i-1,n}\right) \qquad -n < i < n \end{aligned} \tag{3.40}$$

[12]See for example Duffie [27, chapter 3] or Hull [39, chapter 23]

$$q_{n+1,n+1} = \frac{1}{2} q_{nn} d_{nn},$$

where $d_{in} = (1 + r_{in})^{-h_{n+1}}$ is the discount factor at node (i, n). The price at $t = 0$ of a security with payoffs C_{in} at step n of the tree is

$$V_0 = \sum_i q_{in} C_{in}. \tag{3.41}$$

The first input to which the tree is calibrated is a yield curve of discount bonds. The yield curve is used to determine the function $a(t)$ conditional on $b(t)$ and $G(r)$. Discount bonds have payoffs $C_{in} = 1$ at all nodes at their maturity date. For maturity $\tau = t_{n+1}$ the price is

$$\begin{aligned} P_0(\tau) &= \sum_{i=-(n+1)}^{n+1} q_{i,n+1} \\ &= \sum_{i=-n}^{n} q_{in} d_{in} \\ &= \sum_{i=-n}^{n} \frac{q_{in}}{(1 + G^{-1}(a(t_n) + b(t_n) \cdot i/\sqrt{n}))^{h_{n+1}}} \end{aligned} \tag{3.42}$$

The bond price is monotone decreasing in $a(t_n)$. Numerically solving for $a(t_n)$, given $b(t_n)$ and $G(r)$, is a simple univariate problem. The entire sequence $a(t_n)$ is found recursively by moving forwards through the tree.

After calibrating to the yield curve, the resulting tree is used to price a set of swaptions. Let $\tau = t_n$ be a swaption exercise date and let T be the maturity date of the underlying swap. For a call option on the swap rate, called a payer's swaption, the payoffs at nodes (i, n) at the exercise date τ are

$$C_{in}(\tau, T) = [X_{in}(T) - X_0(\tau, T)]^+ \, m \sum_{j=1}^{M} P_{in}(T_j), \tag{3.43}$$

where

$X_{in}(T)$	:	swap rate at node (i,n) for a swap maturing at time T,
$X_0(\tau,T)$	:	forward swap rate at time $t=0$ for a swap starting at τ and maturing at T,
$[x]^+$	:	maximum of x and 0,
m	:	tenor of the swap, defined as the reciprocal of the number of cash flows per year,
M	:	total number of cash flows of the swap,
T_j	:	cash flow dates of the swap, $T_j = \tau + m\cdot j$,
$P_{in}(\tau)$	:	price at node (i,n) of a discount bond with maturity date τ.

The forward swap rate $X_0(\tau,T)$ is a function of the initial yield curve,

$$X_0(t,T) = \frac{P_0(t) - P_0(T)}{m\sum_{j=1}^{M} P_0(T_j)}. \tag{3.44}$$

The actual swap rates at node (i,n) are

$$X_{in}(T) = \frac{1 - P_{in}(T)}{m\sum_{j=1}^{M} P_{in}(T_j)}. \tag{3.45}$$

The future actual swap rates require the entire term structure at node (i,n), which is only available after the function $a(t)$ has been calibrated up to time T. For the valuation of the swaption the tree has to be constructed up to time T and then rolled back to time τ. Once the payoffs $C_{in}(\tau,T)$ have been computed, the value of the swaption is given according to equation 3.41:

$$Q_0(\tau,T) = \sum_{(i,n)} q_{in} C_{in}(\tau,T). \tag{3.46}$$

A put option on the swap rate, called a receiver's swaption, has payoffs for which $[X_{in}(\tau,T) - X_0(T)]^+$ is replaced by $[X_0(T) - X_{in}(\tau,T)]^+$. The difference between the two is $X_{in}(\tau,T) - X_0(T)$ at all nodes (i,n). The difference in value between the put and the call is exactly equal to the value of the forward swap with rate $X_0(T)$. By construction the value of this swap is equal to zero, and therefore the payer's and receiver's swaption have the same value.

3.3.2 Two-factor model

Most of the analysis for the single factor model carries over to the two-factor case. The definition of the spot rates is different and, instead of 3.39, becomes

$$G(r_{ijn}) = a(t_n) + \frac{b_1(t_n)}{\sqrt{n}} \cdot i + \frac{b_2(t_n)\sqrt{2}}{\sqrt{n}} \cdot j, \qquad \begin{array}{l} i = -n, \dots, n, \\ j = -(n-i)/2, \dots, (n-i)/2. \end{array} \tag{3.47}$$

The state prices in the two-factor model satisfy the forward recursion

$$q_{ij,n+1} = \frac{1}{2} q_{i-1,jn} d_{i-1,jn} + \frac{1}{4} q_{i+1,j-1,n} d_{i+1,j-1,n} + \frac{1}{4} q_{i+1,j+1,n} d_{i+1,j+1,n}, \tag{3.48}$$

where we define $q_{ijn} = d_{ijn} = 0$ for nodes (i, j, n) that do not exist. Existing nodes are defined in section 3.2 as $i = -n, \dots, n$ and $j = -(n-i)/2, \dots, (n-i)/2$, both with increments of 2.

The function $a(t_n)$ is calibrated in forward recursion to the yield curve similar to equation 3.42. Calibration to swaptions is completely analogous to 3.43 with future swap rates $X_{ijn}(\tau, T)$, discount prices $P_{ijn}(\tau)$ and option payoffs $C_{ijn}(\tau, T)$ all having the additional j subscript.

The two-factor model is obviously more general than a single factor model but also computationally much more demanding. For the same number of time steps the 2-factor model has $\frac{1}{2}(n+2)$ times as many states as the single factor model. It will converge more quickly to the limiting continuous distribution.

3.3.3 Goodness of fit

Swaptions are valued conditional on the parameters that determine the functions $b(t)$ and $G(r)$. Best fitting parameters are found by minimizing the relative pricing errors between the computed swaption prices and observed swaption prices using least absolute deviations[13],

$$s = \sum_{j=1}^{K} \sum_{\ell=1}^{L} \left| \frac{Q_0(\tau_j, T_\ell)}{Q(\tau_j, T_\ell)} - 1 \right|, \tag{3.49}$$

where

[13] We also worked with squares instead of absolute values, but did not find important differences in results.

τ_j	:	expiration date of swaption,
T_ℓ	:	swap maturity,
$Q(\tau, T)$	:	observed swaption price using Black's formula and the observed implied volatility,
$Q_0(\tau, T)$	:	computed swaption price from the lattice model.

Swaption volatilities are observed for K different expiration dates and L different swap maturities. Since a single factor model will not be able to price all swaptions perfectly, the minimized value for s will not be zero. The least squares criterion is formulated on relative pricing errors, since all swaption prices will generally increase with the time to expiration and the time to maturity. A measure of fit for absolute pricing errors would give a large weight to long term options on long term swaps. Relative pricing errors distribute the weights more evenly.

Implicit in the goodness of fit criterion is that the yield curve is always calibrated exactly using the drift adjustment parameters $a(t_n)$. We experimented with relaxing this constraint by making it part of the objective function. Let $X(T_i)$ be the observed swap yields. The augmented objective function adds an additional term

$$\Psi \sum_{i=1}^{M} |X_0(T_i) - X(T_i)|,$$

to the original objective 3.49, where Ψ is the weight attached to fitting the swap curve relative to fitting swaptions. In that case we also need to parameterize $a(t_n)$, for example using a cubic spline. We found that the relaxation of the yield curve constraint did not lead to much improvement for the swaption prices, because the weight Ψ must be very large. If Ψ is too small the swap curve differs significantly from the observed term structure, implying that forward swaps with rate $X_0(T)$ have non-zero value. As a result, payer's and receiver's swaptions have different values. Because option values are small compared to swap values, relative pricing errors of receiver's swaptions are huge if the interest rate lattice is calibrated on the yield curve and payer's swaptions, even if pricing errors for the latter options are small.

In table 2.3 pricing errors are listed for swaps traded at February 29, 2000 for different maturities, given discount factors resulting from the cubic spline on zero coupon bond prices. As these errors do not significantly differ from zero we use the yield curve from chapter 2 for swap valuation. Note that this pricing error is equal to the difference between an at-the-money payer's swaption and an at-the-money receiver's swaption with the same strike rate and maturity. As by construction payer's swaptions and receiver's swaptions must be equally priced, the yield curve must be exactly matched. We therefore only report results with exact calibration of the yield curve.

3.3.4 Optimizing parameter values

For given parameters $b(t_n)$ and function $G(r)$, the sequence $a(t_n)$ is numerically found by forward recursion, exactly fitting the term structure. Given previous interest rates, the current drift $a(t_n)$ is matched to the corresponding zero-coupon bond price $P_0(t_{n+1})$ according to 3.42. The resulting interest rate lattice is used to value swaptions. Swaption prices are compared to observed prices to obtain the goodness of fit measure s. The optimal values of the parameters $b(t)$, minimizing s, are found numerically using the Broyden, Fletcher, Goldfarb and Shanno (BFGS) method, a quasi-Newton search method.

An outline of the BFGS method is presented here, based on the NEOS Guide [69]. We consider the optimization problem

$$\min_{\beta \in \mathbb{R}^n_+} V(\beta)$$

where the parameter vector β consists of the five coefficients of the smoothed volatility function 3.35, including α and θ. Alternatively, the exponentially decaying function 3.30 might be used as well. The function $V(\beta)$ determines the sum of swaption pricing errors s. Since no closed form expression for $V(\beta)$ exists, the function value can only be evaluated numerically. For given β, the goal function $s = V(\beta)$ is computed as described in the previous section. Our goal is to minimize s by updating β.

For updating β first and second order derivative information is required. The gradient is not calculated exactly (since no closed form for $V(\beta)$ is available), but approximated by considering discrete deviations from the current β variables. The j^{th} element of the gradient of $V(\beta)$ (that is, the derivative of $V(\beta)$ with respect to β_j) is obtained by the approximation

$$\nabla V[j] = \frac{V(\beta + \delta e_j) - V(\beta - \delta e_j)}{2\delta} \qquad \forall j = 0, \ldots, J$$

where e_j is the unit vector with the j^{th} element equal to 1, δ is a small positive number and J is the number of coefficients.

Second order derivative information is captured by the Hessian matrix, which is approximated by using gradient information from past iterations. We initialize the Hessian at the identity matrix. The Hessian improves gradually towards the true second order derivative. Given β^k, the coefficient vector β in iteration k, the corresponding gradient $\nabla V(\beta^k)$ and the approximate Hessian H^k, the linear system

$$H^k d^k = -\nabla V(\beta^k) \tag{3.50}$$

is solved to generate the improving direction d_k. The next iterate is found by setting

$$\beta^{k+1} = \beta^k + \alpha^k d^k \tag{3.51}$$

and defining

$$\begin{aligned} s^k &= \beta^{k+1} - \beta^k \\ y^k &= \nabla V(\beta^{k+1}) - \nabla V(\beta^k), \end{aligned}$$

where α^k is set to the first element of the series $2^{-l}, l = 0, 1, 2, \ldots,$ that satisfies the curvature condition $(y^k)^T s^k > 0$. Now the improved approximation of the Hessian matrix is obtained by the BFGS update

$$H^{k+1} = H^k - \frac{H^k s^k (H^k s^k)^T}{(s^k)^T H^k s^k} + \frac{y^k (y^k)^T}{(y^k)^T s^k},$$

satisfying the quasi-Newton condition

$$H^{k+1} s^k = y^k.$$

Because it takes a lot of computation time to solve a system such as given in equation 3.50, due to inverting Hessian matrices, we can achieve the same result by updating the inverse of the Hessian matrix

$$B^k = (H^k)^{-1}.$$

Now a BFGS update of H^k is equivalent to the following update of B^k:

$$B^{k+1} = (I - \frac{s^k (y^k)^T}{(y^k)^T s^k}) B^k (I - \frac{y^k (s^k)^T}{(y^k)^T s^k}) + \frac{s^k (s^k)^T}{(y^k)^T s^k}.$$

The improving direction can now be directly found by

$$d^k = -B^k \nabla V(\beta^k).$$

FIGURE 3.3: Yield curves for all input dates.

The figure shows the input yield curves for February 29, 2000, February 15, 2001, June 1, 2001 and July 2, 2001, respectively. The time axis displays months, where $T = 0$ is the corresponding date. The yields $y(0, T)$ on the vertical axis are stated in percentages per month.

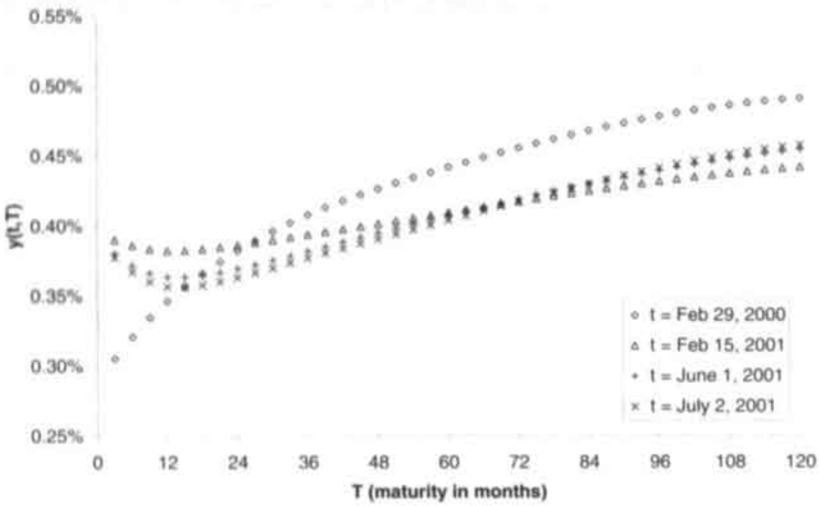

The updated and improved β^{k+1} follows from equation 3.51. Every iteration, $V(\beta)$ is calculated until a (local) minimum is attained. Notice that each iteration requires many calculations to obtain the model prices of all swaptions and a univariate optimization phase to optimize all drift parameters fitting the term structure. An accurate starting solution, formed by the initial vector β, might accelerate convergence towards a minimum pricing error.

3.4 Data

We use data on Euro swap curves for four randomly chosen dates in 2000 and 2001 (February 29, 2000, February 15, 2001, June 1, 2001 and July 2, 2001). The swap data are augmented with EURIBORs having monthly maturities up to 1 year. Both sets of data are pooled to construct a term structure of discount yields using cubic splines. Input term structures for all dates are depicted in figure 3.3. All yield curves are upward sloping. The first one monotonically; the other three with a slight initial inverse hump. Such a negative hump will be extremely difficult to capture for standard equilibrium term structure models like CIR or Vasicek.

For the same dates we also have observations for a matrix of implied volatilities of swaptions with various maturities for the underlying swap and various expiration dates of the option. We consider options with expirations after 1 month, 3 months, 6 months and 1 to 5 years on forward swaps with maturities between 1 and 10 years. We calibrate a lattice for a ten year horizon. Swaptions for which the sum of the expiration date and maturity of the swap exceed ten years are discarded. This leaves 62 swaptions for calibration. Swaption data are available as volatilities implied by Black's model. They are transformed to prices using Black's formula,

$$C_0(\tau,T) = X_0(\tau,T)\left(N(\frac{1}{2}\sigma\sqrt{\tau}) - N(-\frac{1}{2}\sigma\sqrt{\tau})\right)\left(m\sum_{j=1}^{N} P_0(T_j)\right). \tag{3.52}$$

Swaption volatilities for the dates considered are shown in figure 3.4. The volatility surfaces show a more or less pronounced hump. Volatilities are highest up to the middle of the option expiration range. The hump is strongest for options with the shortest swap maturity. The pattern is typical for implied volatilities.[14]

3.5 Results

We first present calibration results and implied interest rate dynamics for the one-factor model. After that we discuss the two-factor model.

3.5.1 One factor model

Table 3.1 presents an overview of the fit, measured by the average relative swaption pricing errors, for three volatility specifications (constant, exponential and square root) and two specifications for the distribution of spot rates (HL and BDT). The volatility specifications are formalized in equations 3.28, 3.33 and 3.35, respectively.

[14]The example in table 22.3 in Hull (2003) exhibits the same pattern for instance.

FIGURE 3.4: Swaption volatilities for all input dates.

The figure displays Black swaption volatilities for all four input instances. The horizontal base is formed by a matrix of option expiration dates and swap maturities (in months). The vertical axis shows implied volatilities in percentages.

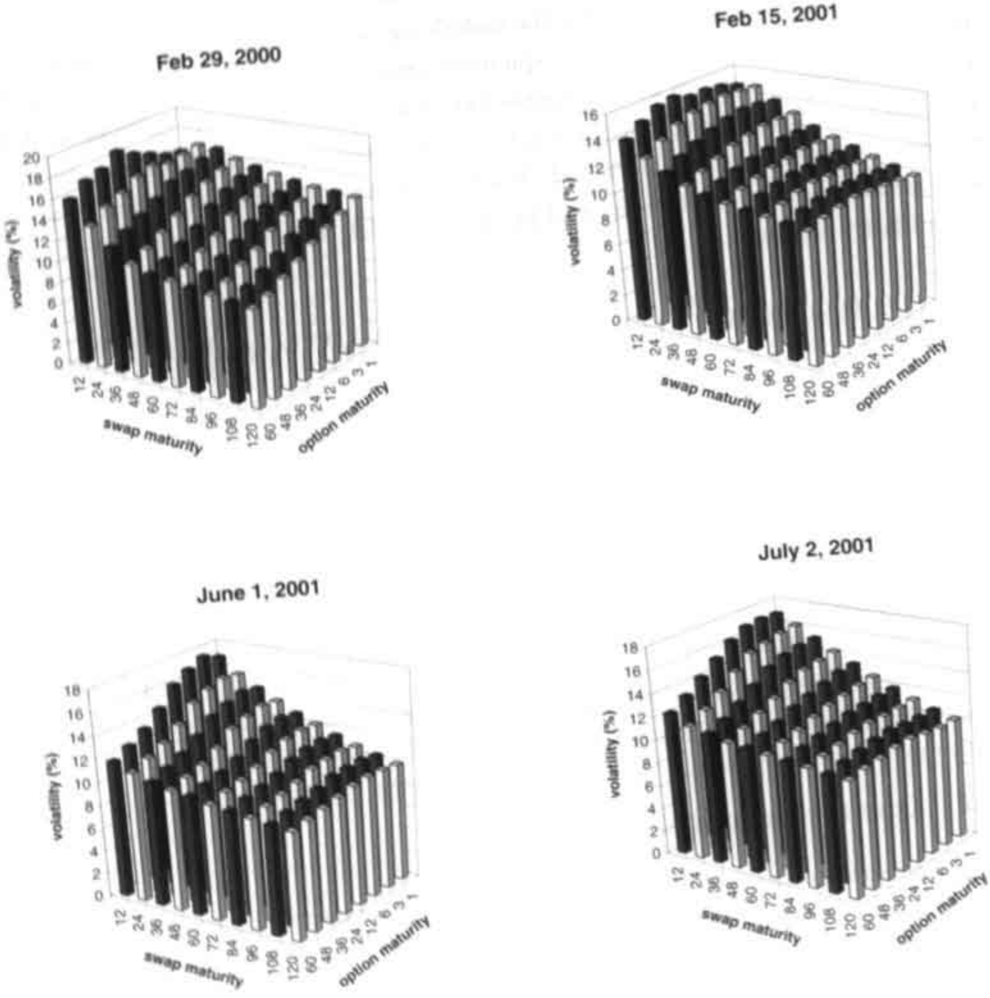

TABLE 3.1: Average swaption pricing errors

The table reports the average pricing error $s/62$ (in percentages, see equation 3.49) over all 62 swaptions with joint time to expiration and maturity of the underlying swap less than or equal to 10 years. N is the number of equally spaced time steps in the ten-year lattice. The "Normal" model uses $G(r) = r$, the "Lognormal" specification has $G(r) = \ln r$. Volatility is either constant (RW) or exponential (EXP) or square root (SQRT). Exact specifications for the volatility functions are in 3.28, 3.33 and 3.35, respectively.

		Lognormal			Normal		
Date	N	RW	SQRT	EXP	RW	SQRT	EXP
Feb 29, 2000	120	10.84	5.87	6.61	9.61	6.21	6.90
	240	13.50	6.26	5.92	8.35	6.71	5.76
	480	12.60	5.65	5.19	7.60	5.73	5.16
Feb 15, 2001	120	11.11	9.92	10.02	11.06	10.55	10.26
	240	9.82	5.97	6.10	7.27	6.95	6.24
	480	9.55	6.31	6.48	7.19	7.18	6.60
June 1, 2001	120	10.97	7.06	7.61	8.27	7.48	7.66
	240	11.60	3.44	3.73	6.11	3.77	3.63
	480	10.95	3.44	3.71	5.61	3.73	3.61
July 2, 2001	120	10.58	7.27	7.42	8.48	7.77	7.66
	240	11.10	4.03	4.29	5.16	4.89	4.16
	480	10.53	4.42	4.34	4.85	4.74	4.19

The lattice horizon equals 10 years. The number of periods in the lattice varies. Computationally it proved very efficient to start by calibrating parameters for monthly periods, that is, $N = 120$. Due to the small number of states and time steps, numerical calibration works fast in such a small lattice. The optimal parameters for $N = 120$ are used as input to redo the calibration with $N = 240$ (or semi-monthly time steps). Since the parameters in the volatility function remain fairly stable with respect to the number of time steps, iterations for finding the optimal parameters for $N = 240$ converge in a few steps when we use the results from $N = 120$ as starting values in the numerical optimization. The same applies a fortiori for the refinement to $N = 480$ ('weekly' time steps) with the solution for $N = 240$ as starting values.

In general pricing errors are smaller the more steps we take. As a crude measure, moving from $N = 120$ to $N = 240$ reduces the average error by about 2%. Further doubling the number of steps to $N = 480$ leads to an additional reduction of only 0.2%. Most of the improvement for the overall fit comes from the short term swaptions. With a single time step per month there would only be two states at the first expiration date. Here the increase from $N = 120$ to $N = 480$ is most important. Since results hardly change after 4 steps a month, we further report the results for weekly time steps only.

The BDT model with constant volatility is the worst performing model. Surprisingly, the results for the HL model with constant volatility are often very close to the models with a more flexible volatility specification. Normal and lognormal models have a similar fit. The BDT model is slightly better in combination with a square root decaying volatility function, whereas the HL model performs a little better when volatilities decrease exponentially. In general, the BDT model with square root decaying volatilities performs best.

Figure 3.5 provides more details on the residuals. We focus on the BDT model for which volatilities decrease according to the square root function. Other model types (excluding the random walk process for volatilities) show similar patterns. Errors are defined as model prices minus observed prices as a percentage of the observed swaption price. For the systematic effect of expiration and maturity on the pricing errors we consider the regression model,

$$v_i = c_0(\text{date}) + c_1 T_i + c_2 \tau_i + \eta_i, \tag{3.53}$$

where v_i is the pricing error for expiration τ_i and swap maturity T_i, and $c_0(\text{date})$ an intercept that varies over the four dates considered. Pooling over all model specifications (excluding the random walk models) the coefficients are $c_1 = 0.11$ and $c_2 = -0.22$. For each option expiration date the longer maturity swaps are overpriced more (or underpriced less), contrary to the results of Driessen, Klaassen and Melenberg [26], but supporting the findings of Fan, Gupta and Ritchken [31]. Short maturity options are usually overpriced, while long maturity options are underpriced by the model, contrary to the results of Longstaff, Santa-Clara and Schwartz [56]. The intercept is negative for all four instances, meaning that short term options on short maturity swaps are underpriced by the model.[15]

Repeating the regression with the absolute errors $|v_i|$ and including the quadratic terms T_i^2 and τ_i^2 we find

$$|\hat{v}_i| = c_0(\text{date}) - 0.13\tau_i + 0.003\tau_i^2 - 0.30T_i + 0.003T_i^2. \tag{3.54}$$

The model fits best for options on five year swaps. Best fitting options are those with maturity around 2 years.

[15] Short swaptions have low prices and are very sensitive to small interest rate changes.

FIGURE 3.5: Swaption pricing errors.

The figure shows relative swaption pricing errors on four different dates, as model price minus observed price divided by observed price. Pricing errors result from the BDT model where volatility parameters decay at rate $\sqrt{t}$. Short term options on long term swaps are mostly overpriced, while long term options on short term swaps are underpriced. The underlying lattice has 480 periods.

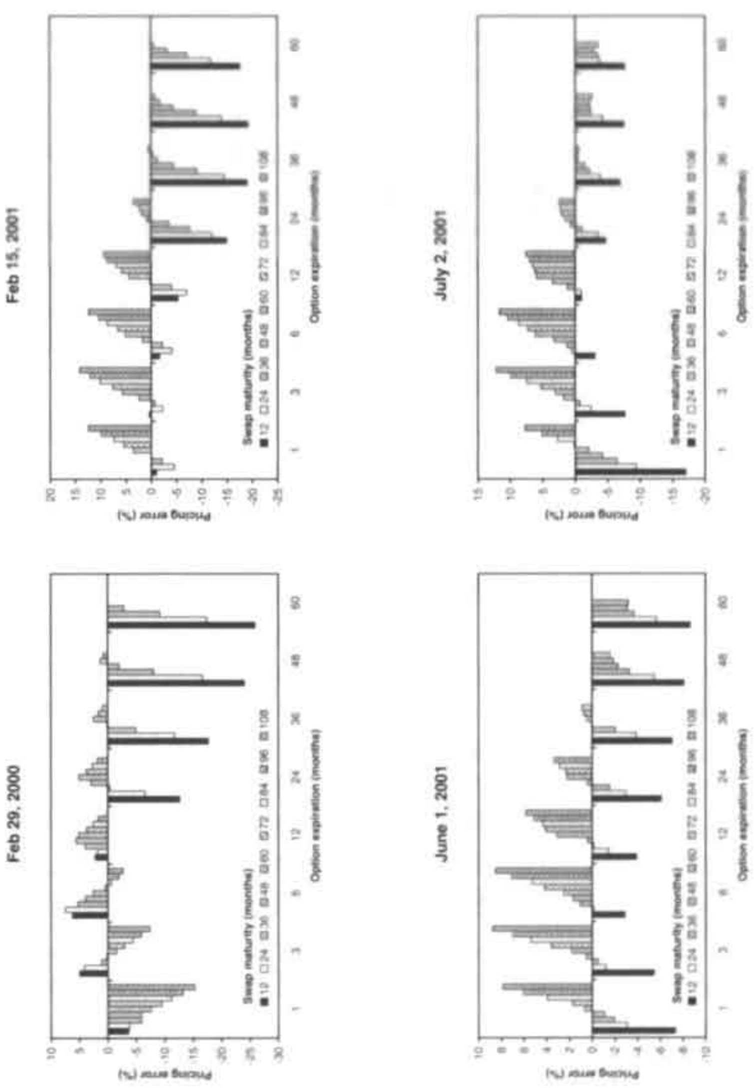

To understand how volatility parameters should change to overcome this pricing error pattern, the following division of nodes could be made. First, divide the time periods in two categories: before option maturity and between option maturity and swap maturity. After swap maturity, volatilities do not affect the swaption price. Second, consider the upper half (representing the higher interest rates, for which payer's swaptions have positive values) and the lower half (for which receiver's swaptions have positive values) of the lattice. This divides the nodes of the lattice in four groups. As a consequence, four channels exist that may influence swaption prices via volatility changes. The first two affect the value of a payer's swaption. This value decreases if the volatility after option maturity decreases or if the volatility before option maturity increases. The first channel states that the resulting decrease in interest rates in the upper half of the lattice implies increasing discount factors and hence the fixed payments at option maturity increase. With higher payments, the present option value decreases. The second channel can be explained by an interest rate increase before option maturity, decreasing discount factors, and therefore payoffs at option maturity are discounted more. As a result the present option value decreases.

Channels three and four affect the receiver's swaption price. The value of a receiver's swaption decreases if the volatility after option maturity decreases or if the volatility before option maturity decreases. Interest rates in the lower part of the lattice are higher (less extreme, closer to the mean short rate) when volatility is small. Higher interest rates after option maturity decrease discount factors, the value of fixed revenues at option maturity decreases. Therefore the current option value decreases. Finally, increasing interest rates before option maturity lead to decreasing discount factors and a lower option value.

Payer's and receiver's swaptions are priced equally. Both payer's and receiver's swaptions with a short option maturity and a long swap maturity are overpriced. Their value can only be reduced simultaneously by decreasing the volatility after option maturity. After all, decreasing volatility *before* option maturity would decrease the receiver's swaption value, but also raise the payer's swaption value. But a reduction in the short rate volatility in later periods would decrease the value of longer maturity options with payments in these periods (for example a 5 year option on a 5 year forward swap) as well. These swaptions are already underpriced. Hence, a tradeoff is found between underpricing and overpricing options. The BDT model, with only one factor to explain the volatility behavior, is not able to improve on this.

An important question might be why this tradeoff in pricing errors exists as described

FIGURE 3.6: Swaption fit February 29, 2000: implied volatility differences.

The figure shows absolute implied volatility differences for February 29, 2000, as model volatility minus observed volatility. The model volatility is the implied volatility of the model swaption price. Errors result from the BDT model with square root smoothing on volatility parameters. We observe the same error pattern as in terms of price differences. The underlying lattice has 480 periods.

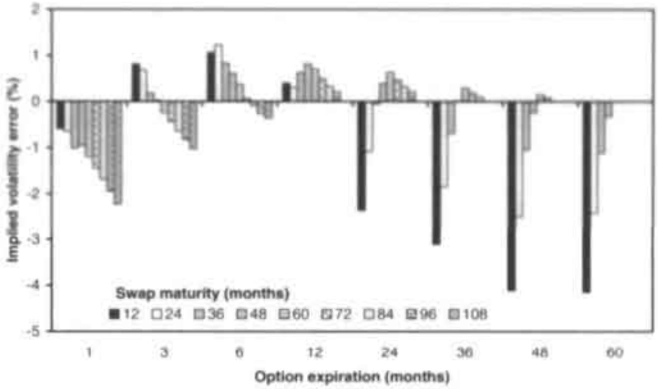

before (and shown by all calibration results) and not opposite (underpricing of short maturity swaptions, especially on long maturity swaps, and overpricing of long maturity swaptions). To overcome this particular pattern of underpricing short maturity swaptions, the volatility after option maturity should increase, also raising the value of already overpriced long maturity swaptions. This effect may be prevented due to the imposed restrictions. The current smoothing function for volatility parameters (decreasing at rate $\sqrt{t}$) implies lower instantaneous short rate volatilities in the future. Consequently, the observed pricing error pattern is more consistent with the volatility specification. However, other smoothing functions (such as exponential decay or even a constant volatility parameter) have the same limitations, perform worse and are theoretically incorrect in this framework.

For the calibration, we minimized relative option pricing errors. A more interpretable (though computationally more cumbersome) measure are implied volatilities. However, the pricing error patterns formed by implied volatilities and relative option prices (as

Table 3.2: Random walk parameters

The table reports the calibrated constant volatility coefficients $\sigma(t) = \sigma$. Entries refer to the monthly volatility of $\ln r$ ("Lognormal") or r ("Normal") and are reported as $100\ \sigma$. N is the number of time steps in the 10 year lattice.

Date	N	Lognormal	Normal
Feb 29, 2000	120	3.876	0.0185091
	240	4.190	0.0199602
	480	4.181	0.0197502
Feb 15, 2001	120	3.402	0.0151339
	240	3.571	0.0156826
	480	3.526	0.0154907
June 1, 2001	120	3.304	0.0149798
	240	3.432	0.0155379
	480	3.404	0.0154444
July 2, 2001	120	3.428	0.0157789
	240	3.561	0.0162809
	480	3.528	0.0161666

in equation 3.49) are very similar. As an example we plot the implied volatilities for February 29, 2000 for the BDT model with the square root volatility function in figure 3.6. The average absolute difference between observed implied volatility and the model's implied volatility equals 0.86% which is of similar magnitude as in Peterson, Stapleton and Subrahmanyam [72].

Parameters of all specifications and input data are reported in tables 3.2 to 3.4. The parameters themselves are not very informative. The differences between the models are better visible by plotting the functions $\sigma(t)$, $\kappa(t)$ and $b(t)$. These functions are shown below.

Beginning with $\sigma(t)$, shown in figure 3.7, we find that all square root specifications are hump shaped. In contrast, only one of the exponentially decaying volatility functions is hump-shaped. All others are monotonically decreasing. All specifications cross the constant random walk volatility approximately in the middle of the ten year horizon of the lattice.

The exponential and square root specifications differ both in the beginning and at the

TABLE 3.3: Parameters square root volatility

The table reports calibrated coefficients for the volatility function

$$\begin{aligned} \sigma(t) &= 0.01 \left(F(t) \cdot \frac{\beta_0}{\sqrt{t}} + (1 - F(t)) \cdot (\beta_1 + \beta_2 t) \right), \\ F(t) &= \frac{\alpha t^\theta}{1 + \alpha t^\theta}, \end{aligned}$$

for both the "Lognormal" and "Normal" distributions for spot rates r. N is the number of time steps in the 10 year lattice.

	Lognormal					Normal				
N	β_0	β_1	β_2	θ	α	β_0	β_1 $\times 10^3$	β_2 $\times 10^3$	θ	α
Feb 29, 2000										
120	12.58	1.012	0.390	1.59	0.017	0.0201	1.980	2.029	1.50	0.015
240	14.50	3.365	0.275	1.47	0.024	0.0202	3.257	2.599	1.49	0.020
480	12.48	2.267	0.394	1.54	0.021	0.0189	5.286	2.090	1.47	0.017
Feb 15, 2001										
120	20.74	0.159	0.117	1.04	0.073	0.0345	6.371	0.979	1.27	0.025
240	16.34	2.161	0.162	0.96	0.131	0.0346	7.385	1.162	1.27	0.030
480	13.29	2.854	0.180	0.98	0.112	0.0346	7.385	1.137	1.27	0.030
June 1, 2001										
120	19.35	0.000	0.186	1.28	0.062	0.0169	7.361	1.134	1.30	0.023
240	14.56	2.649	0.296	1.21	0.095	0.0000	16.896	0.879	1.21	0.027
480	11.49	2.929	0.360	1.21	0.089	0.0000	32.270	1.658	1.19	0.026
July 2, 2001										
120	14.37	0.000	0.343	1.28	0.059	0.0113	9.268	0.943	1.26	0.020
240	9.42	2.639	0.565	1.26	0.089	0.0111	10.616	1.297	1.27	0.028
480	14.50	0.804	0.430	1.28	0.081	0.0109	11.633	1.115	1.26	0.026

TABLE 3.4: Parameters exponential volatility

The table reports calibrated coefficients for the exponential volatility function

$$\sigma(t) = 0.01\,(\theta_0 + \theta_1 t)\,e^{-\kappa t},$$

for both the "Lognormal" and the "Normal" distributions for the spot rates r. N is the number of time steps in the 10 year lattice.

		Lognormal			Normal		
Date	N	θ_0	θ_1	κ	θ_0	θ_1 $\times 10^4$	κ
Feb 29, 2000	120	3.0869	0.1907	0.0223	0.01232	7.813	0.0186
	240	4.7123	0.1332	0.0200	0.01693	6.840	0.0182
	480	4.7039	0.1223	0.0194	0.01576	7.031	0.0182
Feb 15, 2001	120	3.2763	0.0557	0.0107	0.01191	3.309	0.0117
	240	4.5250	0.0264	0.0093	0.01718	1.368	0.0080
	480	4.3944	0.0190	0.0078	0.01659	1.077	0.0066
June 1, 2001	120	4.1740	0.0072	0.0058	0.01359	2.418	0.0100
	240	4.9996	0.0158	0.0102	0.01794	0.650	0.0062
	480	4.8844	0.0077	0.0085	0.01756	0.138	0.0033
July 2, 2001	120	3.4185	0.0898	0.0154	0.01295	2.806	0.0100
	240	4.9299	0.0452	0.0133	0.01747	1.149	0.0068
	480	4.7496	0.0459	0.0129	0.01722	1.024	0.0063

end of the time horizon. The first few months the square root function is increasing steeply. This is offset at the end, where it is much flatter than the exponential.

The differences are more pronounced for the scale function $b(t)$. The scale functions $b(t) = \sigma(t)\sqrt{t}$ are shown in figure 3.8. For a mean reverting process the scale must approach a constant level for large t. By construction the exponentially decaying volatility functions show a decreasing variance of interest rate levels. Whether the theoretical problem is empirically relevant depends on the parameters κ and θ_0/θ_1. For the calibrated parameters in a ten year lattice the difference with the random walk model is clearly visible after a five years horizon. From that point on the random walk still increases at rate $\sqrt{t}$ while the exponential function quickly flattens. For three out of four instances it is already decreasing well before the end of the ten year period. This is the main drawback of the exponential decay of the volatility function.

The square root function implies higher standard deviations for the level of interest

FIGURE 3.7: Instantaneous short rate volatilities.

The figure displays instantaneous volatility curves $\sigma(t)$ for all four dates considered. Each panel shows three different smoothing functions for the volatility parameters: constant (rw), square root (sqrt) and exponential (exp). For the square root functions the hump is more pronounced and the future decline is more moderate than for exponential decay functions. The constant volatility provides an average volatility level. The underlying lattice has 480 periods.

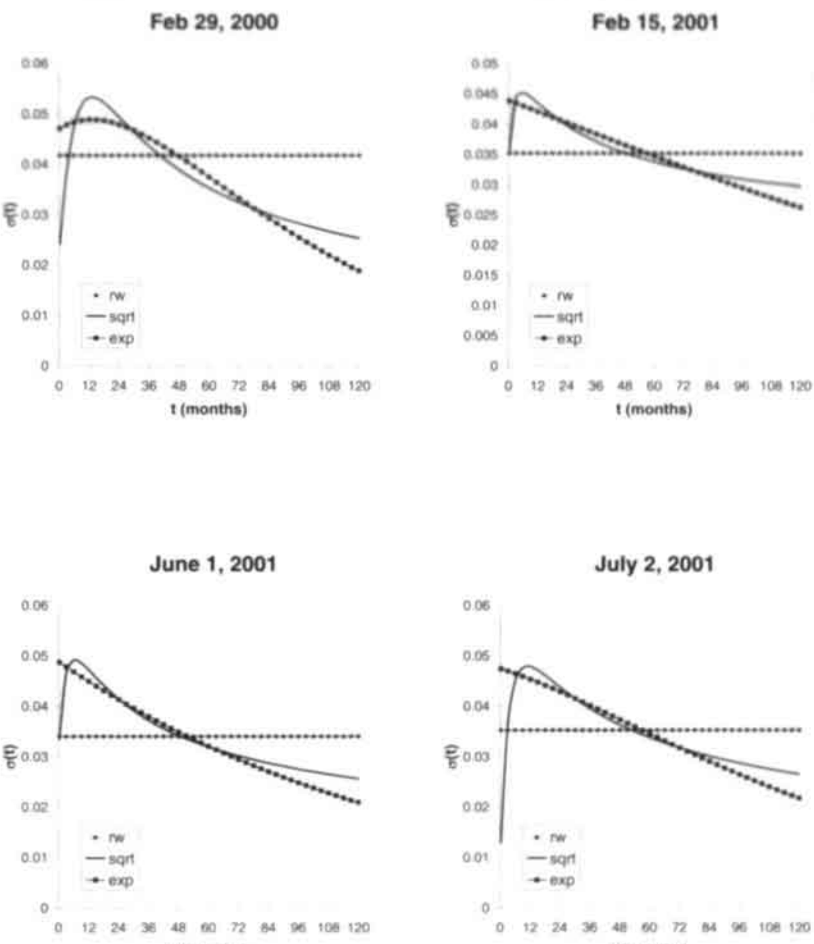

FIGURE 3.8: Scale parameters.

The figure displays scale curves $b(t) = \sigma(t)\sqrt{t}$ for all four dates considered. Each panel shows three different smoothing functions for the scale parameters: constant (rw), square root (sqrt) and exponential (exp). The square root functions approach a constant for large t, while the exponential decay functions imply a decreasing future scale. The underlying lattice has 480 periods.

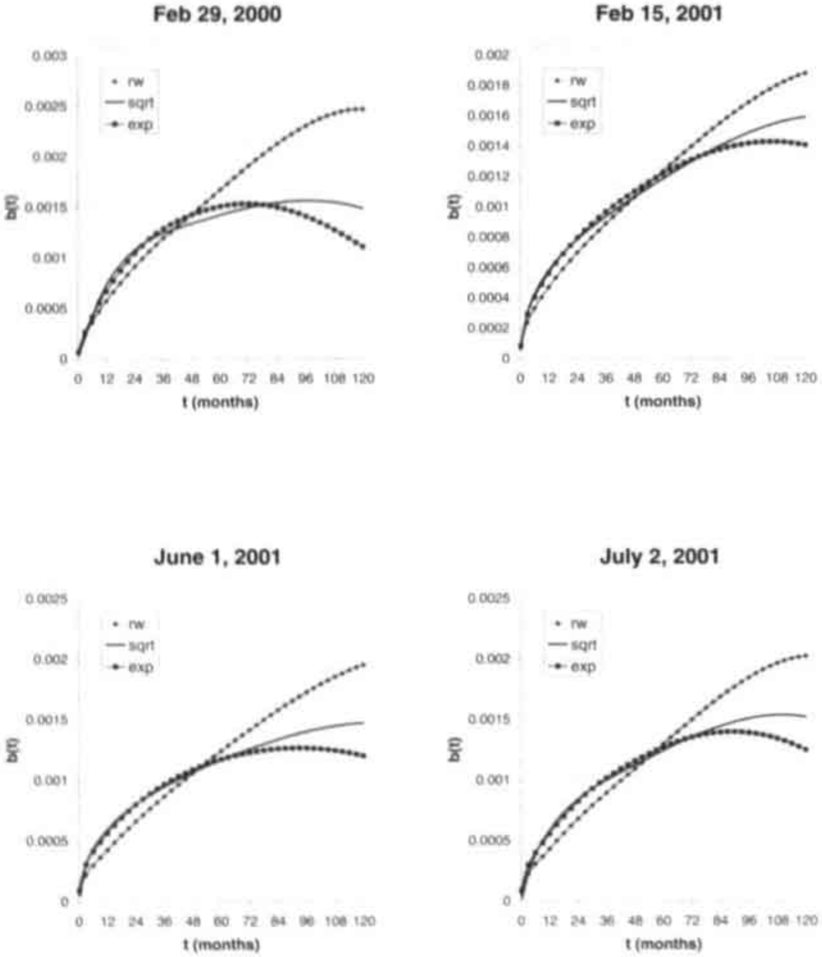

rates $b(t)$ than the exponential model. The increase in variance is at a slower rate than for the random walk. It is remarkable that both the exponential and the square root functions are steeper than $\sqrt{t}$ in the first year (sometimes first two years). This means that initially interest rates display explosive behavior (negative mean reversion). A better description is that within the lattice there is positive autocorrelation for changes in the spot rates. This feature corresponds with the actual time series properties of interest rates.

Interestingly the crucial parameter θ is mostly between one and $1\frac{1}{2}$. The weight function $F(t)$ converges to one at such a slow rate that the scale function $b(t)$ does not converge to a constant, even though $\sigma(t)$ vanishes in the limit. The square root function is constructed such that it satisfies the implications with respect to volatility and scale, approaching a constant level for all data sets.

The mean reversion functions $\kappa(t) = -\sigma'(t)/\sigma(t)$ are depicted in figure 3.9. According to volatility functions decaying at rate $\sqrt{t}$, interest rates diffuse during the first year, after which mean reversion starts.

Figure 3.10 plots the drift parameters $a(t)$ for all instances for the BDT model. Only the curves for the square root volatility function are shown. The results for the constant volatility and exponentially decaying volatility are indistinguishable from the curves shown. The shape of the drift curve is comparable to the shape of the corresponding yield curve. Differences occur because the drift parameters refer to short rates and are median interest rates. This implies that the yield curve is mostly determined by the drift correction and hardly by the dynamic properties of the spot rate process. The dynamics, modelled through the volatility function $\sigma(t)$, only serve to fit the swaption prices.

To gain some insight in the yield curve dynamics, figure 3.11 shows the five future term structures belonging to the states after 1 month ($n = 4$) for the first date (February 29, 2000). The initial yield curve is shown in figure 3.3. Contrary to the implications of most term structure models, the curves are closer together at short maturities than at long maturities. The explanation is that dynamics are not mean reverting initially and mean reversion, if any, is only limited afterwards. The other models and other dates show a similar term structure behavior.

Figure 3.12 shows the final period distributions according to the BDT and the HL model for the square root model on February 29, 2000. Since $\theta > 1.5$ in this instance the distribution converges to a limiting stationary distribution. Both HL and BDT have

FIGURE 3.9: Mean reversion

The figure shows mean reversion rates $\kappa(t) = \frac{\sigma'(t)}{\sigma(t)}$ for BDT for the four dates considered. Volatility is smoothed by exponential and square root functions. Mean reversion only starts after the first year. The amount of mean reversion is very small for large t, while initially a diffusion of interest rates occurs. The underlying lattice has 480 periods.

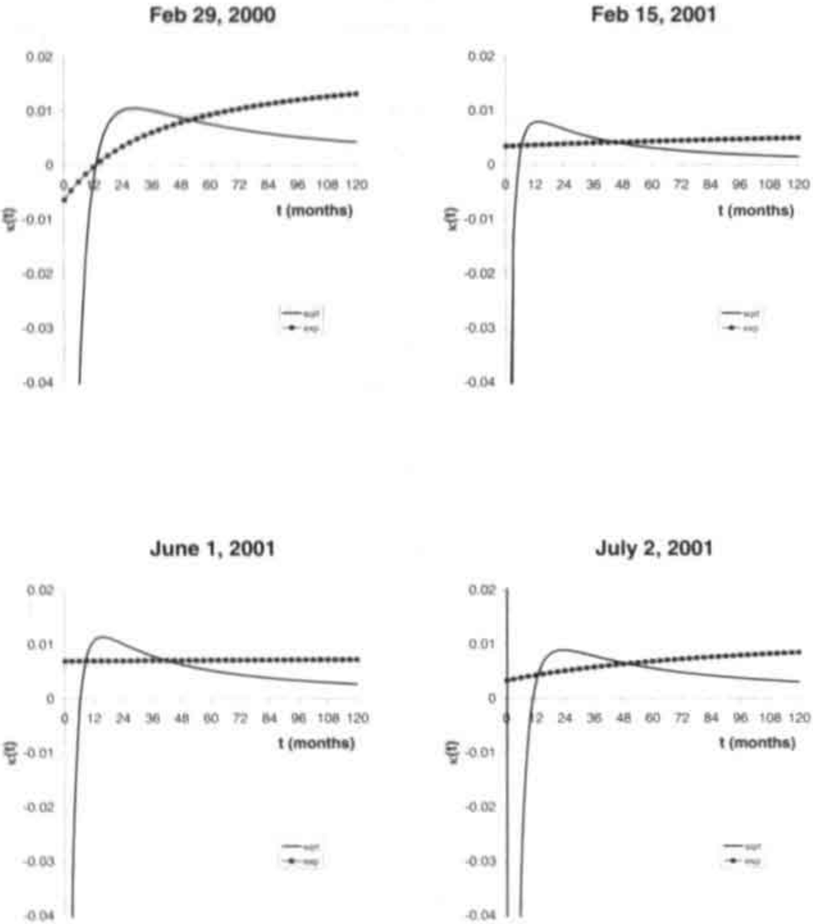

FIGURE 3.10: Drift parameters.

The figure shows the resulting drift parameters $a(t)$ for the four dates considered. Drifts corresponding to the constant and exponential smoothing functions for the BDT model are similar to the square root function. The underlying lattice has 480 periods.

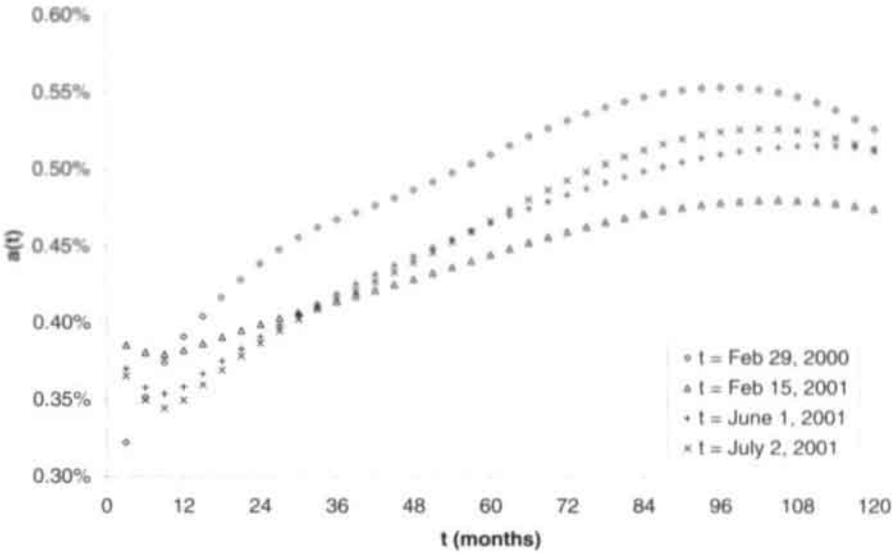

converged after 10 years.[16] Although both models fit the at-the-money swaptions equally well, the difference between the normal and log-normal distributions will be important for the valuation of out-of-the-money options.

[16]Note that all models converge to a stable distribution in the final periods, even Feb 15, 2001 for which $\theta < 1$. Models having $\theta < 1.5$ imply a non-stationary distribution of the short rate. This non-stationarity is not observed in the final periods, but becomes evident when comparing distributions after five and ten years.

FIGURE 3.11: Term structures after 1 month

The figure displays future term structures for the February 29, 2000 instance, starting after 1 month (5 states, as $n = 4$). The underlying lattice has 480 periods. The underlying model is BDT with a square root volatility specification. The term structures diverge initially as interest rates diffuse. For large t the yield curves do converge very little, due to minimal amount of mean reversion. Compare the initial term structure in figure 3.3.

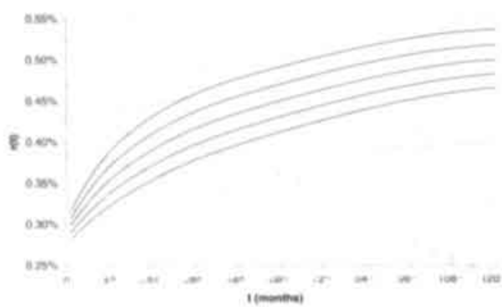

FIGURE 3.12: Short rate distribution February 29, 2000.

The figure shows the short rate distribution for the final period, for BDT and HL models with a square root decaying volatility function. The instance February 29, 2000 is considered. The distribution has converged to a lognormal respectively normal one.

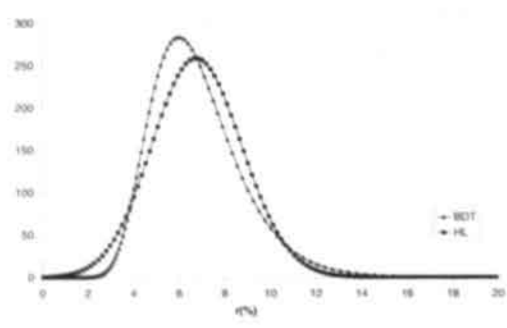

TABLE 3.5: Average swaption pricing errors two-factor model

The table reports the average pricing error $s/62$ in equation 3.49 (in percentages) over all 62 swaptions with joint time to expiration and maturity of the underlying swap less than or equal to 10 years. N is the number of equally spaced time steps in the ten-year lattice. Entries refer to the "Lognormal" model $G(r) = \ln r$ with the square-root specification for both volatility functions in (3.37).

Date	$N = 120$	$N = 240$
Feb 29, 2000	4.27	4.19
Feb 15, 2001	6.59	5.96
June 1, 2001	4.99	3.24
July 2, 2001	5.31	3.72

3.5.2 Two-factor model

For monthly periods ($N = 120$), calibration results improve significantly when considering a two-factor model. The improvement over a one-factor model is at least 25% in terms of average swaption pricing errors. Table 3.5 shows the average pricing errors for the four dates in 2000 and 2001. We only consider the BDT model with square root volatility function as in equation 3.37.

Despite the large improvement due to the second factor, larger improvements have been achieved for the one-factor model when increasing the number of periods from 120 to 240 for three out of four instances. Only for the February 29, 2000 instance an improvement over the best one-factor model solution is obtained with the two-factor model with monthly periods. This observation gives rise to doubling the number of periods of the two-factor model. Table 3.5 also presents the average pricing errors for 240 semi-monthly periods. Computationally, 240 periods in a two-factor model proved to be the maximum.

To enhance convergence to an optimal calibration of the two-factor model with 240 periods, an appropriate initial solution is chosen from known coefficients. Either, the optimal coefficients of the two-factor model with 120 periods or those of the one-factor model with 240 periods are used. Our initial coefficient set is the one that gives lowest average pricing errors. That is, for February 29, 2000 we start with the two-factor model coefficients for monthly periods. For all other instances, the optimization process is initialized by the one-factor model coefficients for semi-monthly periods. The coefficients for the second factor are initialized to small values, as not to disturb the initial optimum too much, but still large enough in order not to get stuck in the local optimum of the one-factor model. Hence, a small initial deterioration of the average pricing error is allowed, in order to find sufficient improvements, leading to a solution that is better than both the one-factor solution with semi-monthly periods and the two-factor solution with monthly periods.

Although all instances are best calibrated by the two-factor model with semi-monthly periods, the difference with the second best model (either the two-factor model with monthly periods or the one-factor model with semi-monthly periods) is never larger than 8%. More detailed calibration results are provided in table 3.6, where all coefficients for the four instances and for monthly and semi-monthly periods are presented. Note that the coefficients θ_i are mostly close to 1.5.

Figure 3.13 displays pricing errors for all swaptions. Model prices are obtained by the BDT model with square root volatility. The underlying lattice has 240 semi-monthly periods. The error patterns are comparable to the one-factor model patterns. From the graphs showing swaption fits we conclude that the short term options are priced closer to the observed price, whereas the long term options have comparable pricing errors to the one-factor model. Especially for the first instance (February 29, 2000) we observe a significant improvement of the pricing errors of the short maturity options.

In the two factor BDT model two sources of volatility exist. Figure 3.14 shows the instantaneous volatilities $\sigma_1(t)$ and $\sigma_2(t)$. Most of them display a hump shaped pattern as in the one-factor model. For some instances, the second factor volatility is much less significant compared to the first one.

For the last two instances, June 1 and July 2, 2001, the scale of the second factor (although very small compared to the first factor) approaches a constant, but the curve is hump shaped. This effect is due to the strictly decreasing volatility curve. Because both β_{20} and β_{22} are nearly zero, the volatility of the second factor is approximately equal to $\beta_{21}(1 - F(t))$, which is an inverse logistic curve.

As for the one-factor model, the drift curve is calibrated exactly to the zero-coupon bonds for all dates in order to price payer's and receiver's swaptions equally. Consequently, the drift curve of the two-factor model closely resembles the drifts of the one-factor model.

TABLE 3.6: Parameters of the two-factor square root volatility model

The table reports calibrated coefficients for the volatility functions

$$\sigma_i(t) = 0.01\left(F_i(t)\cdot\frac{\beta_{i0}}{\sqrt{t}} + (1-F_i(t))\cdot(\beta_{i1}+\beta_{i2}t)\right)$$

with

$$F_i(t) = \frac{\alpha_i t^{\theta_i}}{1+\alpha_i t^{\theta_i}}, \qquad i = 1, 2,$$

for the "Lognormal" model $G(r) = \ln r$. N is the number of time steps in the 10 year lattice.

N	β_{10}	β_{11}	β_{12}	θ_1	α_1	β_{20}	β_{21}	β_{22}	θ_2	α_2
Feb 29, 2000										
120	9.00	1.443	0.375	1.49	0.023	2.11	0.758	0.585	1.50	0.039
240	9.00	1.443	0.391	1.49	0.023	2.11	0.763	0.583	1.50	0.039
Feb 15, 2001										
120	5.62	0.705	0.616	1.62	0.034	0.00	0.000	0.200	1.40	0.007
240	16.14	2.183	0.164	0.96	0.129	0.11	0.033	0.002	1.00	0.049
June 1, 2001										
120	4.33	0.520	0.000	1.48	0.027	9.63	1.533	0.744	1.51	0.027
240	13.95	2.563	0.317	1.23	0.089	0.00	1.110	0.002	1.55	0.034
July 2, 2001										
120	5.43	0.283	0.558	1.49	0.035	15.22	0.000	0.187	1.48	0.024
240	9.98	2.523	0.464	1.24	0.082	0.00	1.942	0.000	1.51	0.032

FIGURE 3.13: Swaption fit two-factor model.

The figure shows relative swaption pricing errors for four dates, as model price minus observed price divided by observed price. Pricing errors result from the BDT model with square root smoothing on volatility parameters. The pricing error pattern is similar to the one-factor model, pricing errors are slightly smaller, especially for short term swaptions. The underlying lattice has 240 periods.

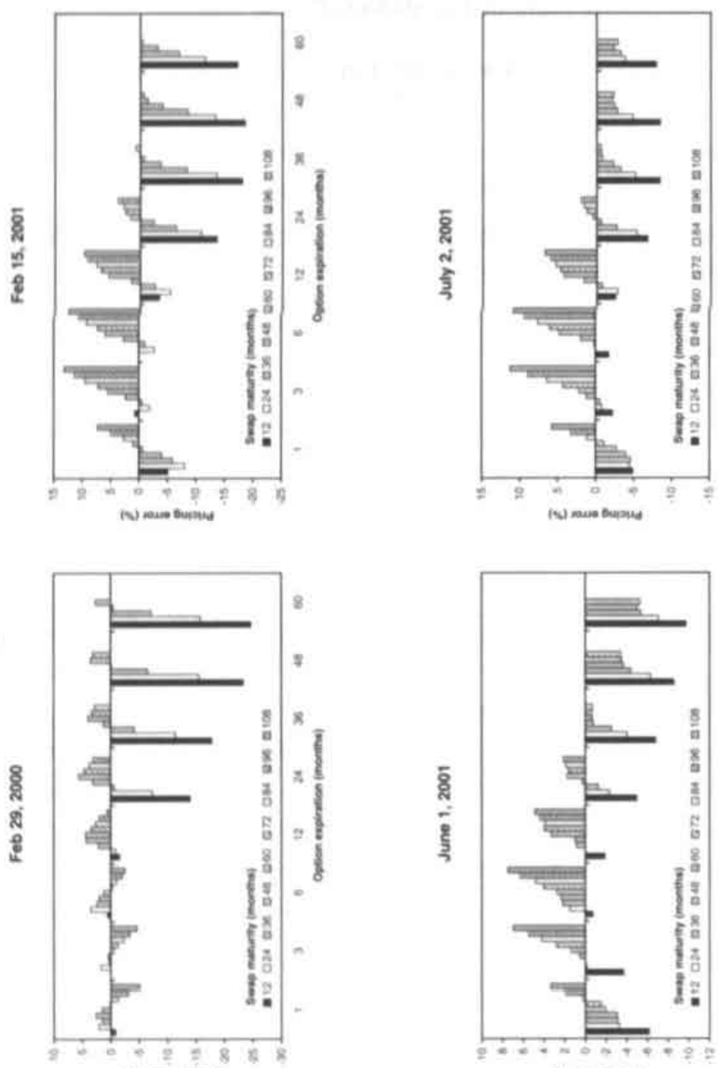

FIGURE 3.14: Instantaneous volatility two-factor model.

The figure displays instantaneous volatility curves $\sigma_i(t)$ for both factors of the two-factor model, for all four dates considered. Results for the BDT model with square root volatilities are shown. The underlying lattice has 240 periods.

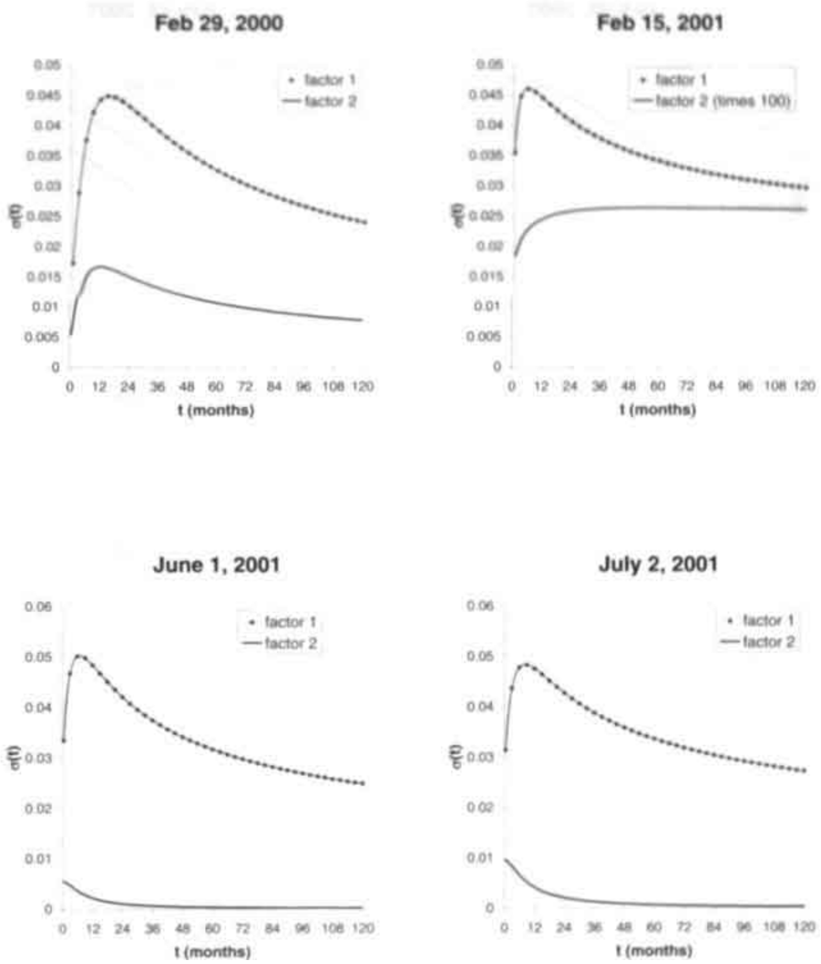

FIGURE 3.15: Scale two-factor model.

The figure displays the scale $b_i(t) = \sigma_i(t)\sqrt{t}$ for both factors of the two-factor model, for all four dates considered. Results for the BDT model with square root volatilities are shown. The underlying lattice has 240 periods.

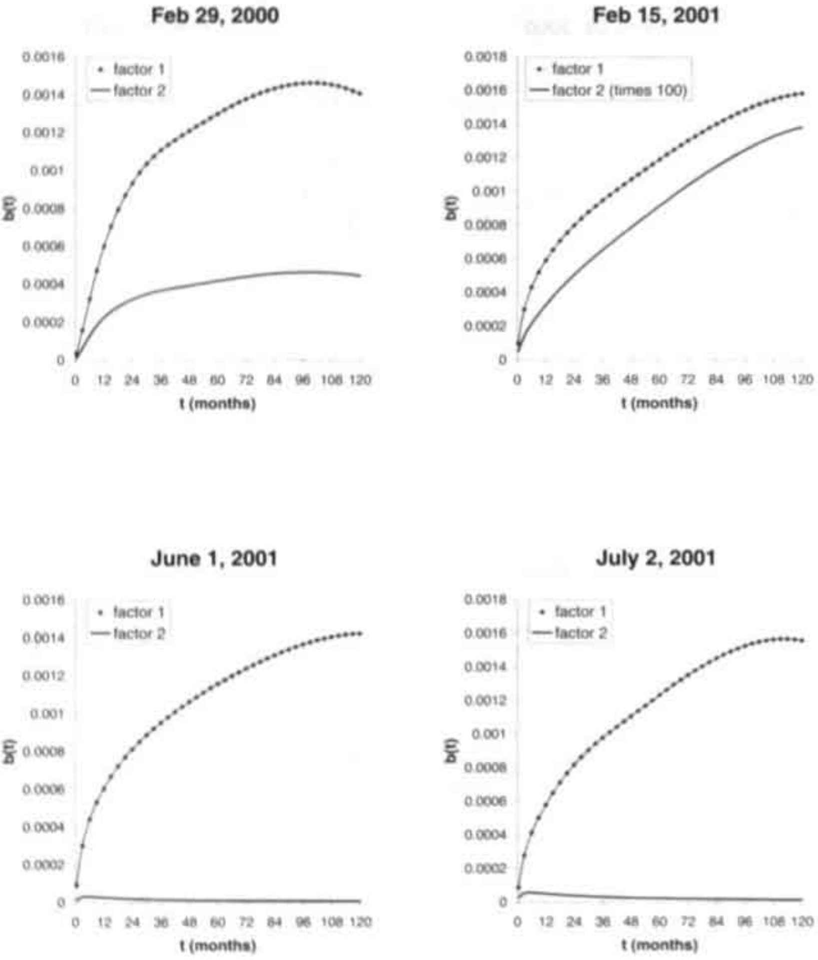

FIGURE 3.16: Mean reversion two-factor model.

The figure displays the mean reversion $\kappa_i(t)$ for both factors of the two-factor model, for all four dates considered. Results for the BDT model with square root volatilities are shown. The underlying lattice has 240 periods.

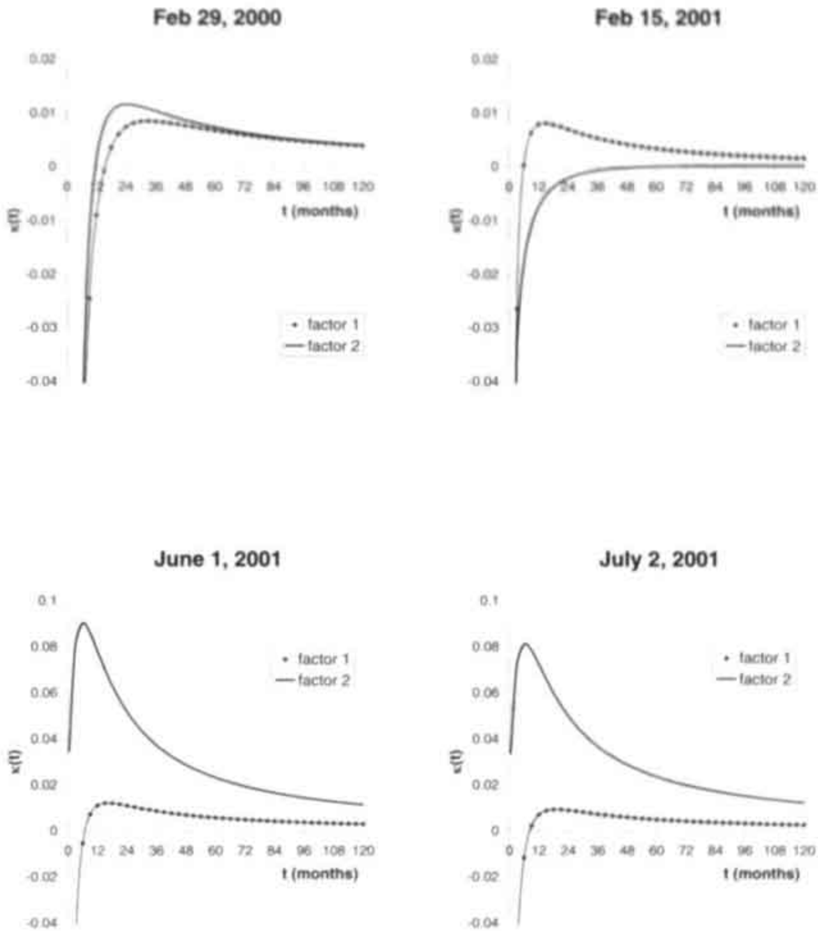

3.6 Conclusion

The calibration performance of a term structure model depends on many interest rate characteristics, most importantly volatility and mean reversion. Applying a one- or two-factor model, based on Ho and Lee [38] or Black, Derman and Toy [9], limits the freedom of parameter choice to describe these characteristics. As shown in equation 3.15, variable time periods are one way to model volatility and mean reversion independently. To have positive mean reversion of interest rates, we show that the length of time steps must be decreasing. However, as long as the implied mean reversion, defined by $\kappa_0(t) = -\sigma'(t)/\sigma(t)$, does not deviate too much from the desired mean reversion, step sizes can remain constant.

Both theoretically and empirically, the volatility parameters $\sigma(t)$ are best smoothed by a function decaying at rate $\sqrt{t}$. In order to have a stationary distribution of the short rate, the smoothing parameter θ must be larger than 1.5. Calibration results suggest that θ is very close to 1.5, as are both coefficients θ_i in the two-factor model. This implies that the scale parameter $\sigma(t)\sqrt{t}$ approaches a constant.

The term structure model is calibrated to match observed swap prices and at-the-money swaption prices. To value short term options, one is tempted to consider many small periods initially and increasing the step length for distant future periods. This leads to decreasing mean reversion (or even diffusion) according to equation 3.27, an increase of the short rate volatility $\sigma(t)$ and an explosion of the yield volatility. To price short term options we therefore considered small, but constant, period lengths.

We apply variations of the one-factor models introduced by HL and BDT for calibration. Interest rates are modelled by means of a recombining binomial scenario tree. Volatility parameters are smoothed according to square root decline, exponential decline, or a constant volatility implying a random walk interest rate process. Our analysis is extended by introducing a two-factor model, based on a trinomial recombining tree, where both factors have independent volatility processes. All models use the drift parameter to fit the yield curve of zero-coupon bonds exactly. If the yield curve is not exactly matched, at-the-money payer's and receiver's swaption prices will be different, while these are quoted equally.

The best performing models when considering four periods per month, BDT with square root volatility and HL with exponential volatility, have average swaption pricing errors of approximately 5% over all instances. When including monthly and semi-monthly periods, the BDT model with square root volatility performs best in nine out of twelve instances.

All instances show a similar pattern of swaption pricing errors. Long term options are underpriced by the model, especially on short term swaps. Short maturity options on long swaps are typically overpriced. An analogous pattern is observed when quoting pricing errors in terms of implied volatilities.

After optimizing the model parameters by a quasi-Newton approach, we tried to improve the resulting local optimum using local search. Local search includes more parameters (drift and non-smoothed volatility for each period), but calibration speed is slower. However, no improvement could be found. When starting with a local search method to optimize the volatility parameters (using all drift parameters to exactly match the yield curve), we did also not obtain a better fit and the error pattern remained.

Volatility according to the square root volatility functions is usually hump shaped, although convergence is rather slow. As a result, the scale parameter is increasing even after 10 years. Mean reversion only starts after one year and remains small afterwards. The short rate distribution after a sufficient number of periods is stable and close to a lognormal (normal) distribution for the BDT (HL) model.

The two-factor model slightly improves calibration results, but the error pattern remains. Short maturity options are priced more accurately because short term flexibility increases with the number of states. Although for monthly periods the implied short rate distribution after one month cannot be modelled accurately, and pricing errors can be larger than for the one-factor model with smaller periods, the two-factor model with semi-monthly periods performs best for all instances considered.

Our purpose is to use the calibrated interest rate lattices for mortgage valuation. Mortgages are particularly long term contracts. As pricing errors for long term assets appear to be small and less subject to discretization issues, the interest rate lattices are applicable for mortgage valuation.

Part II

Mortgage Valuation

Chapter 4

Introduction to Mortgage Valuation

4.1 Introduction

Mortgage contracts have developed from simple bank loans to fixed income securities with various embedded options. A Dutch mortgage contract can have a mix of fixed and adjustable contract rates, can be partially prepaid, may offer borrowers the opportunity to change the fixed rate period of the loan, can have the possibility to lock in minimum rates over a certain time span, or includes a combination of these and additional features. Comparing the offers from various lenders (often intermediaries) and distinguishing between the products on the market is a difficult task in the intransparent mortgage loan market.[1] The objective of the second part of this thesis is to provide a meaningful comparison of the basic contracts offered on the Dutch market. The main complicating factor is the valuation of partially callable mortgage contracts, which will be the topic of the two core chapters of this part.

The best offer would be the contract which has the lowest present value of future cash flows. Since for many mortgage loans the future cash flows depend on future mortgage rates offered by banks and on prepayment decisions made by borrowers, a comparison of different loans requires assumptions about interest rate dynamics and prepayment behavior. In a competitive market one could assume that all loans are priced correctly and that the contract rates reflect the embedded options. The interest to be paid is such that the present value of the future cash flows is equal to the nominal value of the loan. In that

[1]For an overview of the range of mortgage loans on offer in the first half of 2003, we refer to the 'Hypothekengids', a publication (in Dutch) of the 'Vereniging Eigen Huis', an association supporting the interests of homeowners in the Netherlands.

case, the various mortgage contracts will not differ in value, but only in their sensitivity to changes in the term structure. Throughout the second part of this thesis we investigate what magnitude of spreads we can expect for popular embedded options.

Our methodology is based on fair rate computations. Given the current yield curve and information about interest rate volatilities, we define the fair rate as the contract rate that makes the present value of the sum of all cash flows equal to the nominal value. This is the interest rate at which a lender cannot expect to make an abnormal profit on the loan. If the contract rate is higher (lower) than the fair rate, implying a mortgage value higher (lower) than the nominal loan value, the bank makes a profit (loss) on the contract equal to the difference in value. The nominal loan value or principal is scaled to 1 for all mortgage contracts. Fair rate differences between contracts are an indication of the value of embedded options. For each contract we obtain the fair rate that is consistent with the current term structure and volatilities reflected in liquid derivatives such as swaptions.

For some contracts we also observe the contract rate quoted by a lender. The difference compared to the fair rate is an indication of the profitability of the market plus a risk premium for default. Default risk is limited for Dutch mortgages. Full information on the credit status of borrowers is maintained in a central database which is accessible to all major lenders in the market. In addition, many contracts are subject to the 'Nationale Hypotheekgarantie' (NHG) which compensates banks for default losses. Banks quote different rates according to default risk. The lowest rate is for loans secured by NHG, the highest for loan values higher than 75% of the value of the house. In our computations we ignore default risk.

Another important assumption underlying our fair rate calculations is the absence of tax effects. In this chapter we do not consider any client specific aspects affecting the choice for a certain mortgage contract or the mortgage value. Tax effects influence mortgage choice in the Netherlands. Mortgage interest payments are fully tax deductible. As a result mortgage loans are cheap credit. Perpetual mortgages are popular because interest payments remain high. One major effect of tax deductibility is the discouragement of prepayment, since the net contract rate paid is typically lower than the fair rate. When prepaying the mortgage loan while this is not allowed, a prepayment penalty must be paid. This penalty is also tax deductible. The decision to move from a contract with a high interest rate to one with a lower contract rate is not affected by tax effects, as long as both contracts are equally taxed.

A second reason for differences between fair and quoted contract rates could be the interest rate model. For the valuation of embedded options we must assume a dynamic process for spot rates. Although the process is calibrated to the current term structure and to volatilities, many degrees of freedom about distributional characteristics and implied dynamics remain. For robustness of the results we will therefore compute fair rates based on several interest rate models.

We focus on optimal prepayment decisions, driven by interest rate behavior. We only deal with mortgage contracts exposed to interest rate risk and exclude combinations of mortgages with insurances or investments, for which other types of risk play a role. Optimal prepayment of Dutch mortgages is limited to a percentage of the initial loan per calendar year, complicating mortgage valuation significantly. Chapters 5 and 6 introduce valuation procedures for Dutch partially callable mortgages. Chapter 5 deals with interest-only mortgages with partial prepayment options, which can be valued by a lattice approach by keeping track of the number of prepayment options and the number of calendar years remaining. Chapter 6 introduces a linear programming model for general Dutch mortgages.

This chapter provides an introduction to the second part of this thesis. Characteristics of popular contracts will be discussed in the next section. Section 4.3 outlines the valuation methodology for fixed rate mortgages, non-callable and fully callable, which is based on a calibrated binomial lattice. The full prepayment boundary is determined. We provide a valuation method for adjustable rate mortgages based on a dynamic programming approach to create and solve subproblems of the original (path dependent) valuation problem. Section 4.4 identifies valuation problems for Dutch mortgages including partial prepayment options.

4.2 Mortgage characteristics

In this section several mortgage characteristics and options are described concerning the amortization schedule of the mortgage, options to call the mortgage before maturity, and options to adjust the contract rate during the lifetime.

4.2.1 Amortization schedule

The periodical payment of a mortgage is denoted by M_t and depends on the principal U_0 (the nominal loan value, which can be scaled to 1), the lifetime of the mortgage T and the contract rate y. With monthly periods, T is stated in months, y is the monthly contract rate and M_t is a monthly payment. The periodical cash flows are determined as if the contract rate would be valid until maturity. Most contracts, however, have rate adjustments after 5, 7 or 10 years.

The unpaid balance U_t at a future date t can be expressed in terms of the unpaid balance at time $t-1$:

$$U_t = (1+y) \cdot U_{t-1} - M_t. \tag{4.1}$$

The periodical payment consists of an interest payment and an amortization part A_t:

$$M_t = y \cdot U_{t-1} + A_t. \tag{4.2}$$

Boundary conditions imply that U_0 is given (scaled to 1) and $U_T = 0$, such that the mortgage loan is repaid at maturity. Prepayments are defined as additional amortization cash flows O_t. If additional prepayments are allowed, then the periodical payment equals

$$M_t = y \cdot U_{t-1} + A_t + O_t. \tag{4.3}$$

An additional prepayment affects the future amortization schedule. All contracts considered in this thesis have the restriction $O_t \geq 0$. Some modern contracts allow for negative O_t as long as $U_t \leq U_0$ for all periods t.

Annuity mortgages

An annuity mortgage is characterized by a constant periodical payment M, such that $U_T = 0$. The periodical payment of a non-callable annuity is given by[2]

$$M = U_0 \cdot \frac{y}{1-(1+y)^{-T}}. \tag{4.4}$$

This follows easily from the fact that the principal equals the sum of all future payments discounted at the contract rate:

$$U_0 = \sum_{t=1}^{T} \frac{M}{(1+y)^t}. \tag{4.5}$$

[2]Based on the *remaining* loan value U_{t-1} the next periodical payment can be obtained by $M_t = U_{t-1} \frac{y}{1-(1+y)^{-(T-t+1)}}$, being equal to M for all t for constant y.

The periodical payment consists of redemption and interest payments. When time evolves interest payments decrease and the periodical redemption amount increases. At date t the interest payment can be shown to be

$$I_t = M \cdot \left[1 - \frac{1}{(1+y)^{T-t+1}}\right], \tag{4.6}$$

whereas the redemption amount equals

$$A_t = M \cdot \frac{1}{(1+y)^{T-t+1}}. \tag{4.7}$$

The unpaid balance of an annuity mortgage is given by

$$U_t = (1+y) \cdot U_{t-1} - M. \tag{4.8}$$

Linear mortgages

A linear mortgage is characterized by a constant redemption amount. The interest payments, and hence the total periodical payments, decrease over time. The periodical redemption $A_t = A$ is given by

$$A = \frac{U_0}{T}, \tag{4.9}$$

where U_0 and T are again the principal and the lifetime of the mortgage. Now the total periodical payment at date t, M_t, of a non-callable linear mortgage is given by the sum of redemption and interest:

$$M_t = A + y \cdot U_{t-1}, \tag{4.10}$$

with y representing the contract rate as before. Finally, the unpaid balance at date t, U_t, is given by the principal minus all redemption payments up to and including period t:

$$U_t = [U_0 - A \cdot (t-1)] \cdot (1+y) - M_t = U_0 - A \cdot t. \tag{4.11}$$

Interest-only mortgages

Interest-only mortgages are free of regular redemption ($A_t = 0$) during the lifetime of the contract. Only interest has to be paid, amortization takes place at the end ($A_T = U_0$).

The final redemption could be done by taking out a new loan, selling the house or using private savings. Without prepayments the unpaid balance remains at the initial level:

$$U_t = U_{t-1} = \ldots = U_0. \tag{4.12}$$

Each period t, interest is paid equal to $y \cdot U_{t-1}$. This is also the periodical payment M_t because of the absence of regular redemption amounts. Concerning the contractual payments, an interest-only mortgage is a long term coupon bond for which the principal is redeemed at maturity.

Since in the Netherlands interest payments on a mortgage loan are fully tax deductible, interest-only mortgages are popular. The unpaid balance does not decline because no redemption takes place. Therefore interest payments and tax deductions remain high.

Savings and investment mortgages

During the lifetime of a savings or investment mortgage only interest payments occur. Redemption does not take place before the end of the mortgage contract. Without prepayments the unpaid balance remains unchanged until maturity. To repay the loan an account is opened at the beginning of the contract. Such an account usually takes the form of a savings or investment fund. At the end of the contract the mortgage is repaid using the account balance.

If we assume that payments to the account are made according to an annuity schedule, a savings or investment mortgage is a combination of an annuity and an interest-only mortgage. Tax deductible interest payments remain high, while periodic redemption payments are used to increase the account and not for amortization purposes. The savings account earns the same interest as must be paid on the mortgage loan. Consolidating the loan and the savings account, the pre-tax cash flows are identical to an annuity mortgage. On an after-tax basis they are different to the client.

Similar to an interest-only mortgage, excluding additional prepayments the unpaid balance remains at the initial level U_0. The savings account B_t increases by the interest earned and by the monthly redemption, redemption being equal to the total monthly payment minus interest:

$$B_t = (1 + y) \cdot B_{t-1} + M - y \cdot U_0. \tag{4.13}$$

The net outstanding loan, denoted by $N_t = U_t - B_t$, can easily be shown to be

$$N_t = (1 + y) \cdot N_{t-1} - M, \tag{4.14}$$

which is similar to the unpaid balance of an amortizing mortgage.

An investment mortgage is similar to a savings mortgage, but is exposed to higher risk since it is uncertain whether the return on the investment fund is sufficient to repay the loan at maturity. Comparable to an interest-only mortgage, savings and investment mortgages are long term coupon bonds for which the principal is redeemed at maturity.

4.2.2 Call options

When entering a non-callable mortgage contract, prepayment of (any part of) the loan is not allowed. Only regular redemption and interest payments occur. Valuation of a non-callable, fixed rate mortgage boils down to a simple present value calculation using the current term structure. It does not require an interest rate model. Non-callable mortgages are hardly issued by banks. A typical mortgage contract includes some form of prepayment opportunity.

A fully callable mortgage can be completely prepaid in a single period. This is a typical feature for US mortgage contracts. Besides the regular periodical payment the mortgagor must decide whether to prepay or not in each period. Interest rate driven prepayment takes place when interest rates are low. Calling the total mortgage loan can be financed by taking out a new mortgage loan at the prevailing lower contract rate.

In the Netherlands, only prepayment of a percentage (10% to 20%, depending on contract and lender) of the principal mortgage loan is allowed every calendar year. If the mortgagor wants to prepay more than this percentage a penalty must be paid equal to the present value of the difference between the future monthly payments of the new contract and the existing mortgage. Because of this prepayment penalty a larger than allowed prepayment is never optimal. A penalty is not paid in case of moving or death of the mortgagor or at a contract rate adjustment date.

For Dutch mortgage contracts, as an alternative to full prepayment, many loans include a 'time for reconsideration' option. A 'time for reconsideration' option allows a borrower to fix the contract rate achieved during the reconsideration period for the next fixed rate period. For many contracts the specified period equals the last year or two years of the current fixed rate period. Clients can choose when to fix the contract rate during this period and will do so when it is observed to be low.

A mortgage with an embedded 'time for reconsideration' option is equivalent to a time constrained fully callable mortgage. The option allows for refinancing the mortgage loan during the specified period, at a future (lower) contract rate, equivalent to calling the complete loan and initiating a new contract at the future rate. Consider a mortgage contract with an option to fix the contract rate to the future mortgage rate in any desirable month t during the last year of the fixed rate period. This is equivalent to calling the mortgage (without penalty costs) in month t and initializing a new contract at the mortgage rate valid at t. Calling is restricted to the last year of the fixed rate period and the new loan must be initiated at the same bank.

Sometimes the lowest interest rate during the first year (or two years) can be fixed as contract rate. This rate is called an entering rate. During an entering period, the mortgage loan is equivalent to a variable rate mortgage. The contract is fair priced since each period the contract rate is adjusted to the one-period fair rate. After the contract rate is fixed, the mortgage loan is comparable to a fixed rate mortgage. For this reason, in the sequel a mortgage including an entering rate option is not considered as a separate mortgage type.

4.2.3 Contract rate adjustment

With a fixed rate mortgage (FRM) an adjustment of the contract rate during the lifetime of the mortgage is not possible, neither for the bank nor for the client. A single contract rate is faced during the total lifetime. Fixed rate mortgages are rare in the Netherlands.

An adjustable rate mortgage (ARM) is defined as a mortgage which contract rate is reset periodically. Usually the time between reset dates is constant. During a period between two reset dates, called a fixed rate period, the contract rate is fixed. Typical fixed rate periods span 5 or 10 years, whereas a common mortgage lifetime equals 30 years. After each fixed rate period the contract rate is reset, based on the future term structure. At every adjustment date full prepayment is allowed without penalty. However, as long as fair rate valuation is considered, the adjusted contract rate is fair and prepayment is not profitable at a reset date. Valuation of ARM's is the topic of section 4.3.2. For a detailed discussion on ARM valuation, see Kau, Keenan, Muller and Epperson [48] and Van Bussel [17].

The contract rate adjustments can be unrestricted or restricted. In case of unrestricted adjustments the future contract rates are solely based on the future term structure and are independent of the initial contract rate. Restricted adjustments might impose caps, floors or both on the future contract rate. A cap (limiting a contract rate increase compared to the current rate) protects the client from extremely increasing mortgage rates, while a floor (limiting a contract rate decrease) protects the bank from decreasing rates. Typical Dutch mortgages do not have adjustment caps or floors.

In the Netherlands many variations on contract rate adjustments exist, two of which are discussed in more detail: a bandwidth mortgage (a variable contract rate which is not adjusted if the contract rate remains within a bandwidth) and an interest rate limit mortgage (combining a fixed rate and a variable rate mortgage by the use of a contract rate limit).

A *bandwidth mortgage* is a variable rate mortgage for which every period a contract rate adjustment may occur. If an adjustment takes place, the new rate is set to the contract rate for a mortgage with the same maturity date (that is, a shorter lifetime). The contract rate is only adjusted if the new rate lies outside a bandwidth. Let us consider a bandwidth on an increase of the contract rate. If the new contract rate is lower than the initial rate, the contract rate is adapted. But if the new contract rate is larger, adjustment only happens when the difference between the two contract rates exceeds the bandwidth. In that case, the new contract rate is set to the higher contract rate minus the bandwidth.

As an example of a bandwidth mortgage, suppose the initial contract rate equals 6%, and the bandwidth, usually depending on the length of the fixed rate period, is 1.25%. If in the next period the newly determined contract rate is 5%, then the contract rate is adjusted to 5%. But if the future contract rate rises to 7% no adjustment takes place. Only if the contract rate rises above 7.25% the rate is changed. For example, if the new rate becomes 8% only 6.75% is paid as the bandwidth is deducted. Many mortgage contracts exhibit an equally sized bandwidth for decreasing contract rates. In that case contract rate decreases are passed on only if they exceed the bandwidth.

A bandwidth mortgage can be interpreted as an adjustable rate mortgage with every period being an adjustment date. Adjustments are restricted and depend on the original rate since the contract rate is only adjusted if the contract rate difference exceeds the bandwidth.

The *interest rate limit* mortgage combines a variable interest rate with a fixed interest rate. Initially the client chooses a fixed rate period and an interest rate limit to be paid at most. First, the variable rate is paid until the contract rate for the chosen fixed rate period exceeds the interest rate limit. From that date on the fixed rate period starts and the fixed interest rate corresponding to the chosen fixed rate period is paid, until this fixed rate period ends.

When paying variable interest rates, the periodical payments of an interest rate limit mortgage can be determined directly. Rate adjustments occur every period and unrestricted, which is comparable to a variable rate mortgage. After fixing the interest rate, the mortgage price can be determined similar to a fixed rate mortgage. The fair rate is equal to the interest rate limit by construction. Consequently, the valuation of interest rate limit mortgages does not require separate treatment.

4.3 Valuation

Binomial tree methods for American mortgage valuation are developed by among others Kau, Keenan, Muller and Epperson [48, 50], Van Bussel [17], and, for multiple state variables, Hilliard, Kau and Slawson [37]. Also (exotic) option literature (see Hull [39] for an overview) is valuable for mortgages with prepayment options.

The underlying process of mortgage valuation problems is captured by an interest rate model. We calibrated a binomial scenario tree of interest rates on market prices of swaps and swaptions to reflect the term structure of interest rates and volatilities. For calibration purposes the scenario tree is allowed to be recombining, and is therefore called a binomial lattice. When valuing mortgages however, we cannot always rely on a lattice approach. For many contracts (for instance ARM's and mortgages with partial prepayments) we must apply a non-recombining tree approach, because the history of prepayment decisions or adjustments is relevant for the current mortgage value, thereby introducing path dependencies. Hence a unique path is required to lead to each node, which is typical for a non-recombining tree.

Non-recombining trees are computationally inefficient compared to lattices, as the number of states increases exponentially. To avoid working with non-recombining trees, algorithms can be developed for decomposing the original mortgage valuation problems into smaller problems, each of which can be solved using a lattice. After all, we prefer to use multiple lattices instead of a single tree. In chapter 5 we will analyze the decomposition of a mortgage contract into several coupon bonds. In chapter 6 we discuss a decomposition method based on the past prepayment strategy. An ARM can be decomposed in line with this last approach, based on past adjustments. This will be the topic of section 4.3.2. First, we analyze the valuation of fixed rate mortgages.

4.3.1 Fixed rate mortgage valuation

In this section we describe valuation methods for both non-callable and callable fixed rate mortgages. These mortgage types are path independent, that is, mortgage prices do not depend on past prepayment decisions, and can be based on a binomial lattice approach. The periods in a binomial lattice are denoted by $t = 0, \ldots, T$. For a period t the number of states equals $t+1$, labelled $i = 0, \ldots, t$. Transitions are defined by two possible successors, $(i+1, t+1)$ and $(i, t+1)$, of each node (i, t).[3] Both successor nodes are reached with probability one half. Node $(0, 0)$ is called the root node.

As analyzed in chapter 3, the state process can be extended to a two-factor binomial lattice (i, j, t), but the costs in terms of computation increase rapidly. To study the robustness of the underlying term structure model with respect to the number of factors included, we will consider both one- and two-factor models for pricing mortgage contracts.

Non-callable mortgage valuation

For the valuation of non-callable fixed rate mortgages an interest rate model is not required. We can use a simple pricing method based on separating a mortgage into zero-coupon bonds with different maturities. The price of a non-callable fixed-rate mortgage can be obtained by discounting the periodical payments using the current term structure:

$$P(0,0) = \sum_{t=1}^{T} \frac{M_t}{(1 + y(0,t))^t}, \tag{4.15}$$

where $y(0,t)$ is the current t-period spot rate. For an annuity the periodical payment is constant: $M_t = M$. Mortgages with different amortization schedules have varying periodical payments M_t.

Although a lattice approach is not required for pricing non-callable mortgages, we introduce the valuation procedure based on a binomial lattice for a non-callable mortgage, before proceeding to fully callable mortgages. During a fixed rate period the contract rate is constant. For a particular rate y the monthly payment M_t is calculated by equation 4.4

[3] Although the tree structure is similar as in the previous chapter, nodes are labelled $i = 0, \ldots, t$. Previously, we used indices $j = -t, \ldots, t$ with increments of 2 and $j = 2i - t$.

if we consider an annuity or by 4.10 in case of a linear mortgage. The unpaid balance at the end of the fixed rate period is found recursively by equation 4.8 or 4.11, respectively.

At the end of a fixed rate period (time τ) either the mortgage is fully amortized or it can be prepaid or refinanced costlessly. Full amortization occurs if τ equals the lifetime T of the mortgage, in case the final unpaid balance $U_\tau = 0$. If the fixed rate period is shorter than the mortgage lifetime ($\tau < T$), a positive unpaid balance remains at time τ, which can be prepaid without penalty. Consequently this remaining unpaid balance can be valued at the current maturity τ discount rate.

Since full prepayment is allowed, the price of a mortgage loan at τ in all states of the world is equal to the unpaid balance: $P(i,\tau) = U_\tau$, $i = 0,\ldots,\tau$. After determining the final mortgage values, the intermediate values $P(i,t)$ of the mortgage are found by working backwards, starting at $t = \tau - 1$, until the current value $P(0,0)$ is obtained:

$$P(i,t) = \frac{0.5 \cdot P(i+1,t+1) + 0.5 \cdot P(i,t+1) + M_{t+1}}{1 + r_{it}}, \quad i = 0,\ldots,t. \tag{4.16}$$

The mortgage price equals the discounted sum of next period's expected price and the periodical payment.

The fair contract rate, that is, the contract rate that results in a mortgage price equal to the principal of 1, can be solved for numerically by increasing (decreasing) the initial contract rate when $P(0,0)$ is smaller (larger) than 1. The fair rate is iteratively determined by a straightforward bisection method. The monotonicity of the mortgage price with respect to the contract rate follows from the fact that a higher contract rate implies an increase in interest payments. Consequently, the price, being the sum of all (discounted) cash flows, rises.

Up to this point the effect of commission costs is excluded. Commission costs, typically 1% of the mortgage loan, can be easily accounted for in fair rate calculations. Including a 1% commission the contract rate is fair if the resulting mortgage price equals 0.99 instead of 1.

A (binomial) lattice approach for non-callable fixed rate contracts provides information about the contract rate risk. The two possible fair rates at $t = 1$, y_u and y_d, are a measure of the current fair rate volatility according to

$$\sigma_y = 0.5 \cdot (y_u - y_d). \tag{4.17}$$

Note that different measures of contract rate risk can be applied, for instance based on the change in the mortgage price.

Fully callable mortgage valuation

Basic non-callable mortgages can be valued by a term structure of interest rates. When pricing fully callable mortgages a lattice approach is required, since the mortgage value does not only depend on time but also on the state of the world. Prepayment typically occurs at low interest rates, while for high interest rates the mortgage loan is continued. Full prepayment options are included in all American type mortgages. Compared to non-callable mortgages, the lattice based valuation procedure for fully callable loans differs in just one aspect. In every period and state of the world, the mortgagor must decide whether to prepay or not. If no prepayment takes place, valuation continues exactly as described by the backward recursion 4.16. If the loan is prepaid, the true mortgage value must equal the value of the remaining unpaid balance. We refer to Kau, Keenan, Muller and Epperson [49] for a detailed discussion on fully callable mortgages and to Hull [39] for American option valuation in general.

In the mortgage valuation lattice, full prepayment will occur if interest rates are low. Following a scenario path, prepayment is triggered as soon as the interest rate drops below a certain level. This level is indicated by the *full prepayment boundary*. We define the full prepayment boundary as the distinction between nodes for which full prepayment is optimal (nodes with the lower interest rates) and nodes for which full prepayment is not optimal (at the higher interest rates).

The regular periodical payment M_t of a fully callable mortgage is obtained by equation 4.4 for an annuity and by 4.10 for a linear mortgage. Unpaid balances are determined according to equation 4.8 or 4.11, as if no prepayment is allowed.

To determine the optimal prepayment strategy, the profitability of calling the mortgage must be evaluated in every node of the mortgage valuation lattice. The value of a fully callable mortgage in node (i,t) without calling is given by

$$P'(i,t) = \frac{0.5 \cdot P(i+1,t+1) + 0.5 \cdot P(i,t+1) + M_{t+1}}{1 + r_{it}}, \quad i = 0,\ldots,t. \tag{4.18}$$

For each node, the mortgage is called if and only if $U_t < P'(i,t)$ during the backward recursion process. The mortgage value in node (i,t) after the prepayment decision is taken equals

$$P(i,t) = \min\Big(P'(i,t), U_t\Big). \tag{4.19}$$

The decision whether to prepay in state (i,t) is based on the optimal prepayment strategy

FIGURE 4.1: Full prepayment boundary

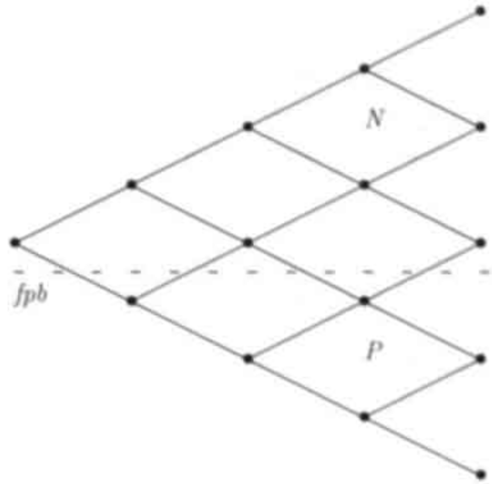

after time t, since the future mortgage price has been adapted if prepayment were optimal in a successor node of (i,t). Now, $P(i,t) = U_t$ if $U_t < P'(i,t)$ (that is, prepayment is optimal) and (i,t) is marked before the recursion continues. When the recursion is completed the marked nodes are separated from the unmarked nodes by the full prepayment boundary. The current price $P(0,0)$ of a full prepayment mortgage is lower then the price of a similar non-prepayment mortgage, because full prepayment reduces the mortgage value. Equivalently, the fair rate is larger to compensate for prepayment advantages.

Calling penalties can easily be accounted for by either multiplying U_t with $(1+c)$ where c is a relative penalty, or by adding C to U_t where C is an absolute penalty. Then prepayment takes place if and only if $U_t \cdot (1+c) < P'(i,t)$ or $U_t + C < P'(i,t)$. Taking the relative prepayment penalty as an example, the mortgage price in node (i,t) observed by the borrower equals

$$P(i,t) = \min\Big(P'(i,t), U_t \cdot (1+c)\Big). \tag{4.20}$$

For positive prepayment penalties the current price of such contract will be larger than for a penalty-free fully callable mortgage, but smaller (or equal) than the price of a non-callable mortgage.

Figure 4.1 shows a small lattice example with a possible full prepayment boundary (the dashed line labelled *fpb*). Note that in general the boundary does not have to be horizontal. Also, a larger interest rate decrease might be required to trigger full prepayment behavior, especially when penalty or commission costs are involved. In the nodes below the boundary (labelled P), interest rates are low and full prepayment is optimal, whereas in nodes above the boundary (labelled N), interest rates are high and no full prepayment is optimal.

We conclude this section with a description of the mortgage valuation algorithm for fully callable mortgages:

- Fix the initial contract rate. Calculate the periodical payment according to equation 4.4 or 4.10. For interest-only mortgages, the periodical payment equals $y \cdot U_0$.
- Calculate the unpaid balance U_τ at the end of the fixed rate period according to equation 4.8, 4.11 or 4.12.
- For all states at the end of the fixed rate period, set the mortgage value equal to the unpaid balance: $P(i, \tau) = U(\tau)$.
- Perform a backward recursion to obtain the continuation value of the mortgage according to equation 4.18. A prepayment decision must be taken in each node. The mortgage price including prepayment decision is given by equation 4.19.
- In case we aim to find the fair rate, increase (decrease) the contract rate if the current mortgage value is smaller (larger) than the nominal loan value (scaled to 1). Repeat the process until the mortgage price equals 1.

4.3.2 Adjustable rate mortgage valuation

After analyzing the valuation of mortgages with a single fixed rate period, we will now deal with adjustable rate mortgages (ARM's), having multiple fixed rate periods. Each fixed rate period has a (fair) contract rate. The current mortgage price can be determined solely based on the contract rate of the first fixed rate period, if and only if the future contract rates are not restricted to the first.

Adjustable rate mortgages have been discussed by Kau, Keenan, Muller and Epperson [48] and Van Bussel [17]. An ARM is a relatively simple mortgage, but one that cannot be valued by a single lattice approach in case the future contract rates are related to the initial contract rate. An ARM is a path dependent instrument (the contract rate in a certain

FIGURE 4.2: Decomposition for ARM

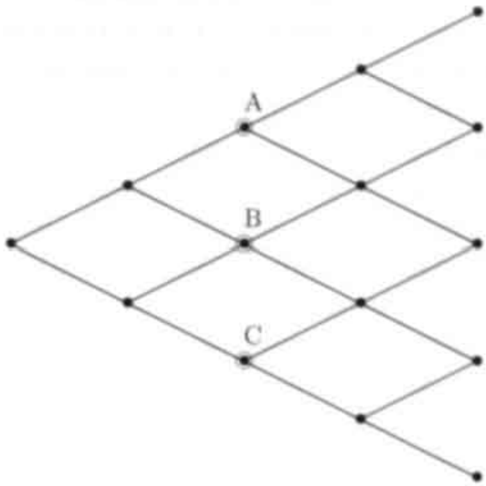

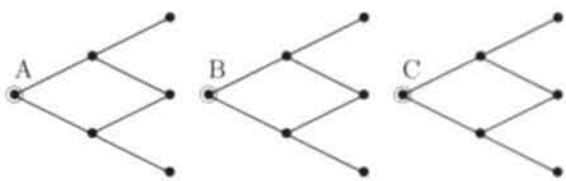

state depends on which adjustment node has been passed). Path dependent contracts are basically valued by a non-recombining tree approach. For efficiency reasons we prefer to decompose the original ARM valuation problem into smaller subproblems, in order to use binomial lattices.

An ARM may be considered as a combination of several mortgages, each spanning a single fixed rate period. Mortgages spanning a single fixed rate period can be priced similar to an FRM. The main idea of the decomposition method is to split up the original ARM valuation problem into several fixed rate periods. In this section we discuss ARM's with possibly several adjustment dates. At each adjustment date the new contract rate becomes the fair rate for the next fixed rate period, if the adaptation to any future fair rate is allowed. Adjustment restrictions such as caps and floors, limiting the allowed increase and decrease of the contract rate respectively, raise the complexity of the valuation of an ARM.

To show the decomposition process, consider a four period lattice as in the upper diagram of figure 4.2. Suppose the first fixed rate period lasts two time steps. After two steps the contract rate is adjusted to the future fair rate, which can have three different values. Assume that the mortgage is fully amortized after four steps, at the end of the lattice. To determine all fair rates during the contract lifetime, we consider the fixed rate periods separately. See figure 4.2 for the decomposition at $t = 2$. The future fixed rate periods can be viewed as fixed rate mortgages with a two-period lifetime. For the root nodes of all sublattices, the unpaid balance is scaled to 1. The periodical payment can be calculated based on a two-period lifetime. Analogous to a usual FRM, the mortgage price is calculated by backward recursion through each sublattice from $t = 4$ to $t = 2$. To obtain the fair contract rates for all future sublattices the process is run iteratively.

When determining the fair rate of the original ARM (for the initial fixed rate period in the original lattice), the same procedure can be used. For a given contract rate, the periodical payment is calculated for a four-period mortgage. Then the unpaid balance at the end of the fixed rate period ($t = 2$) is derived. The possible mortgage prices at $t = 2$ are based on this unpaid balance. If all future contract rates are fair, the mortgage prices $P(i, 2) = U_2$, $i = 0, 1, 2$. Given $P(i, 2)$, backward recursion 4.18 is applied to determine the current price $P(0, 0)$. In case of a fully callable mortgage, intermediate prices are compared to the remaining loan U_t.

Without caps or floors the mortgage value in points A, B and C equals the remaining loan value. Therefore it is sufficient to consider only the first fixed rate period to obtain the initial fair contract rate or the mortgage price. Adjustable rate mortgages including caps and floors may not have feasible fair contract rates for all fixed rate periods. When adjustments are restricted, the initial contract rate affects the range of future rates. For fairly pricing an ARM including caps or floors we must therefore consider all fixed rate periods.

Cap and floor restrictions might imply that fair future contract rates are not feasible, although a fair initial contract rate can still be achieved. Returning to our small example in figure 4.2, let $U_i(0)$ and $P_i(0, 0)$, $i = 0, 1, 2$, be the unpaid balances and prices of an FRM starting at $t = 2$ in the roots of the three sublattices. The initial unpaid balance of each sublattice $U_i(0)$ has been scaled to 1, the corresponding price $P_i(0, 0)$ is therefore based on an initial loan of 1. However, the unpaid balance (of the original four-period mortgage) after 2 periods equals $U(2)$. The correct price in the original lattice is given by $P(i, 2) = P_i(0, 0) \cdot U(2)/U_i(0)$.

If the adjusted contract rate is again fair, $P_i(0,0) = U_i(0) = 1$. In that case $P(i,2) = U(2)$. If a future fair rate is prohibited by cap or floor restrictions, $P_i(0,0) \neq U_i(0) = 1$, resulting in $P(i,2) \neq U(2)$. Concluding, future contract rates affect $P(i,2)$ for some states i, and thereby also the current mortgage value $P(0,0)$ and the initial fair contract rate.

In this section we discussed a decomposition method to value adjustable rate mortgages. Unrestricted fair rate valuation boils down to pricing several fixed rate mortgages. To obtain the initial fair contract rate, considering an FRM maturing at the first adjustment date is sufficient. To calculate all future fair contract rates, all sublattices (corresponding to future fixed rate periods) must be considered. Also, when rate adjustments are restricted all sublattices must be priced. We summarize the pricing algorithm for a fully callable ARM, based on Kau, Keenan, Muller and Epperson [48] and Van Bussel [17]:

- Decompose the ARM valuation problem into FRM problems corresponding to all possible future fixed rate periods. First price the most distant future FRM's and proceed recursively to the current fixed rate period.

- For each fixed rate period, fix the initial contract rate. Calculate the periodical payment according to equation 4.4 or 4.10. For interest-only mortgages, the periodical payment equals $y \cdot U_0$.

- Calculate the unpaid balance U_τ at the end of the fixed rate period according to equation 4.8, 4.11 or 4.12.

- For all states at the end of a fixed rate period, set the mortgage value equal to the unpaid balance. Scale if the fair rate cannot be achieved due to cap or floor restrictions: $P(i,\tau_j) = P_{ij}(0,0) \cdot U(\tau_j)/U_{ij}(0)\ \forall i$ where τ_j is the date of adjustment j and $P_{ij}(0,0)$ and $U_{ij}(0)$ represent the root price and unpaid balance of the corresponding sublattice, initialized in state i at adjustment date τ_j.

- Perform a backward recursion to obtain the continuation value of the mortgage according to equation 4.18. A prepayment decision must be taken in each node. The mortgage price including prepayment decision is given by equation 4.19.

- In case we aim to find the initial fair rate, increase (decrease) the initial contract rate if the current mortgage value is smaller (larger) than the nominal loan value (scaled to 1). Repeat the process until the mortgage price equals 1.

4.4 Concluding remarks

In this chapter we analyzed the main concepts of mortgage valuation using binomial lattices. Non-callable and fully callable mortgages have been covered, as well as fixed and adjustable rate mortgages. Resulting fair rates will be provided and analyzed in chapter 7, in which also fair rates of partially callable mortgages are included.

Because all mortgage contracts issued in the Netherlands are partially callable, the following two chapters deal with the valuation of such contracts. Valuation of mortgages including partial prepayment options is more complicated than the valuation issues discussed in this chapter. Partial prepayments introduce path dependencies. The mortgage price in node (i, t) does not only depend on the prepayment decision in node (i, t), but also on past prepayment decisions, that is, the scenario path towards (i, t). Moreover, prepayment is limited per calendar year, introducing dependencies between prepayment decisions in the same calendar year. A trivial lattice approach is therefore not applicable for the valuation of partially callable mortgages.

In chapter 5 we introduce a valuation approach to circumvent the use of inefficient non-recombining trees. A lattice based backward recursion approach, storing at each node the number of remaining prepayment opportunities and the number of calendar years left proves sufficient to value partially callable interest-only mortgages.

Chapter 6 introduces a linear programming model for the valuation of general partially callable mortgage contracts. This model is based on a non-recombining tree, but LP duality enables us to bound the fair contract rate efficiently and accurately.

Chapter 7 combines fair rates for all mortgage contracts to determine the values of full and partial prepayment options for various mortgage types. Also, the effect of commission costs is analyzed. Several model specifications are applied to improve robustness.

Chapter 5

Optimal Prepayment of Dutch Partial Prepayment Mortgages[1]

5.1 Introduction

The Dutch mortgage market is one of the largest in Europe. Charlier and Van Bussel [19] show that the Netherlands is ranked second among the countries of the European Union in terms of both the outstanding mortgage debt as a percentage of GDP and the outstanding mortgage debt per capita. The size of the total mortgage pool is huge, taking into account the limited number of inhabitants. Increasingly popular are interest-only and savings mortgages. These mortgages are not amortized during the contract lifetime, such that interest payments remain high. They are popular because the Netherlands is the only (European) country where interest payments are fully tax deductible.

An important difference between Dutch mortgage contracts and US mortgages are the prepayment restrictions. Whereas in the US the mortgage can be fully prepaid at the discretion of the borrower, a Dutch contract has limits imposed on the maximum prepayment per year. Only 10 to 20 percent of the initial loan can be prepaid per calendar year.

Empirical prepayment behavior for Dutch mortgages has been documented extensively in studies by Charlier and Van Bussel [19], Hayre [32], Van Bussel [17] and Alink [1]. Optimal exercise of the Dutch prepayment option has not been considered so far, contrary to the large literature on optimal prepayment for US mortgages.[2]

[1]This chapter is based on joint work with P. Schotman.

[2]See e.g. McConnell and Singh [62], Kau et al. [48]-[51] and textbook references such as Hull [39].

The technical problem in pricing a mortgage under the optimal prepayment policy is the path dependence. Prepayment is an American style option that is usually valued by backward induction. With partial prepayments, backward induction is not directly applicable since at the terminal nodes of the lattice it is not known how often prepayment has taken place in earlier years. This chapter derives the optimal prepayment strategy of Dutch interest-only mortgages within a binomial lattice.

Mortgage prices are compared based on fair rates. The fair rate is defined as the contract rate at which the mortgage price is equal to the nominal loan value. Using the optimal prepayment policy, we calculate the typical fair rate spreads between the mortgage rates for a non-callable mortgage and a mortgage with partial prepayments. This spread provides the value of the partial prepayment option and will be compared with the American option value of full prepayment, provided by the contract rate spread between a non-callable and a fully callable mortgage.

5.2 Formulation

We consider the following contract specifications. The contract has a maturity of L years. The contract interest rate ("mortgage rate") is Y percent per year. Contractual payments have the form of a coupon bond. Each period a fraction $y = Y/K$ of the outstanding loan is paid, with K the number of periods per year. No regular amortization takes place. The remaining principal, if any, is repaid at maturity. Before maturity, the principal can be repaid according to the following two conditions:

- The mortgage rate is fixed for $M \leq L$ years, after which it is adjusted to prevailing market rates. No caps or floors apply to the interest rate adjustment. At the adjustment date the borrower has the right of full prepayment.
- In each calendar year a fraction $1/N$ of the principal can be prepaid. The total loan can be prepaid over a period of N years.

Full prepayment is also allowed when the borrower dies or when the house is sold. We will not include the latter effects. We consider optimal call policies that minimize the present value of the cash flows paid by the borrower and received by the lender. The optimal call policy provides a lower bound on the loan value to the lender. For several reasons borrowers might not all follow the optimal prepayment policy. Foremost is the tax incentive. Interest rate payments on a mortgage on a main residence are fully tax deductible, whereas earned interest is not taxed in many cases (for example when household wealth is below a threshold). Another reason is the link between a mortgage loan and a life insurance contract, present in part of the outstanding mortgage contracts.

Since the interest rate is adjusted at the end of year M to the new market rate, and full prepayment is allowed at that moment, the value of the contract at the end of year M will be equal to the principal. Hence, for valuation we do not need to look beyond year M. The contract is in effect a coupon bond with maturity M that is callable in N steps.

The valuation problem is complicated by the path dependence of the partial prepayment option. As an example, consider a 3-year bond that allows prepayment of 50% of the principal each calendar year. For an American option we would like to apply backward induction within an interest rate lattice. But at any node of the lattice in year 3 the value of the contract is path dependent. Its value will differ depending on whether prepayment has taken place in years 1 and 2, or not. The terminal conditions of year 3 will depend on the number of prepayments in previous years.

Path dependence can often be solved by introducing an additional state variable, which keeps track of all possible values of a function of the state variable at each node. This is the technique applied to price American lookback options presented in Hull [39]. The present problem is different. The path dependence is not through a function of the state variable but through endogenous decisions of the borrower. In that sense the partial prepayment problem is related to shout options (see Cheuk and Vorst [21]).

Path dependence does not occur for all partial prepayment mortgages. When $M \leq N$, the loan cannot be fully prepaid until maturity. We can decompose this loan as a portfolio of a noncallable coupon bond, with interest payments being coupons, and a bond with maturity M that is callable in exactly M annual steps. In symbols,

$$V(M,N) = \frac{M}{N}V(M,M) + \frac{N-M}{N}P(M), \qquad M < N, \tag{5.1}$$

where $V(M,N)$ is the value of a mortgage contract with maturity M years that can be called in N annual steps and $P(M)$ is the value of a non-callable coupon bond with the same maturity. If the total principal equals 1, a fraction of $1 - M/N$ is invested in the coupon bond. This part of the loan cannot be prepaid before maturity. A fraction M/N is invested in a bond of which a fraction $1/M$ can be called in each calendar year. To keep the notation as light as possible, we suppress the dependence on the contract rate Y. Valuation of a non-callable coupon bond is trivial. The callable bond, with value $V(M,M)$

can be further decomposed as a portfolio of M simple callable bonds,

$$V(M,M) = \frac{1}{M}\sum_{\ell=1}^{M} C(\ell), \tag{5.2}$$

where $C(\ell)$ is the value of a coupon bond with maturity M that is only callable in year ℓ and not in any other year.

The possibility of partially calling a coupon bond can be excluded, since optimally each callable bond is either prepaid fully or not at all. After partially calling a coupon bond a portfolio remains, which is a linear combination of the continuation value of the coupon bond and the call value. Therefore, partial prepayment is dominated by either no prepayment or full prepayment. If any prepayment is profitable, then full prepayment is most profitable, else no prepayment is optimal. Partial prepayment may only be considered if an investor is indifferent between calling the coupon bond or not, but in that case one of the extreme calling options is just as profitable.

We conclude that options having $M \leq N$ are simple American style options, without path dependence, that can be easily valued by a lattice method. The more complicated contracts are the ones with $M > N$. These are contracts that can be fully prepaid before the contract rate adjustment date. We thus concentrate on pricing mortgage contracts with $M > N$.

The idea for the valuation of a partially callable mortgage is the construction of two lattices. The first lattice has annual time steps and describes the evolution of the prepayments. Figure 5.1 depicts the lattice for a mortgage with a fixed rate period of 7 years and an annual prepayment of 20% of the principal. The original mortgage contract is indicated as the root node (7,5) in the figure. If the borrower has not exercised the prepayment option in the first year, the contract will become a mortgage with maturity of six years of which 20% can be prepaid annually. In the lattice this is node (6,5). The alternative is that the borrower will have exercised the first prepayment option, in which case a mortgage with maturity $M = 6$ remains that can be prepaid in $N = 4$ steps; this is node (6,4).

If we would know the values of the two contracts at the end of year 1, we would be able to price the mortgage contract. For this we can use a regular binomial interest rate lattice with nodes (i,t) with associated discount factors $d_{it} = (1+r_{it})^{-1}$ $(t = 0,\ldots,T_1,\ i = 0,\ldots,t)$. The end of year ℓ occurs at time $t = T_\ell$. Each node (i,t) has two successors, $(i+1,t+1)$ and $(i,t+1)$. Both successor nodes are reached with probability one half.

FIGURE 5.1: Prepayment conversion lattice

The figure shows the evolution over time (in calendar years) with respect to the number of years m and the number of prepayment options n remaining. Nodes are labelled (m, n). Leaves have either the possibility of prepayment in each calendar year, or a single prepayment possibility.

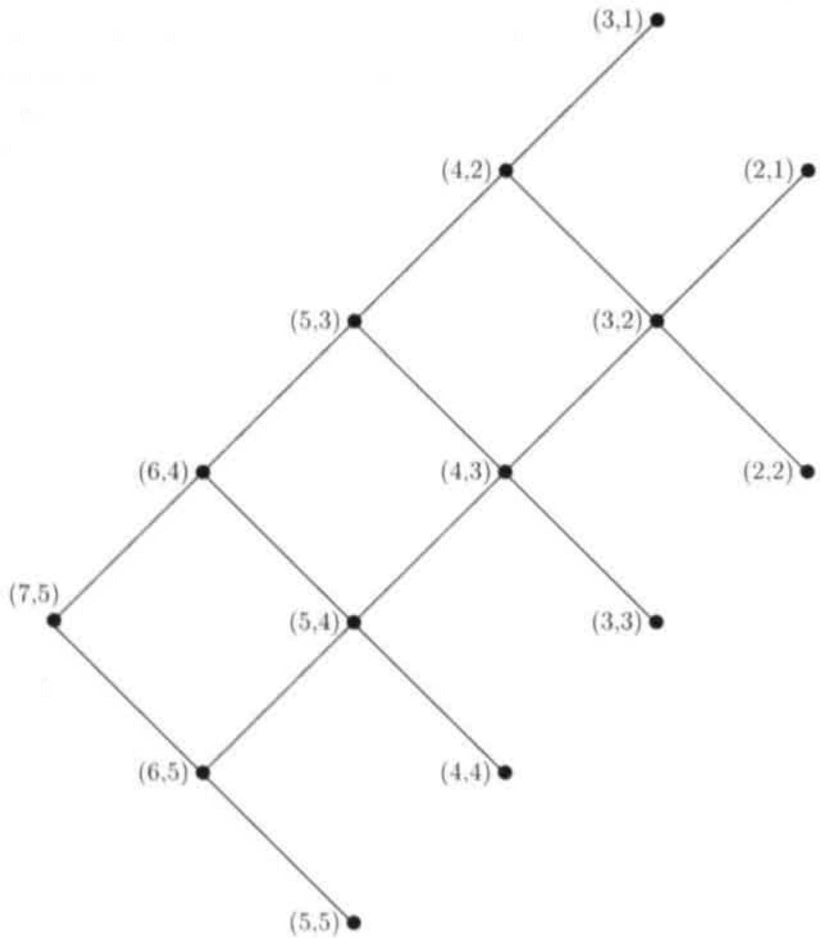

Suppose $V_{i,T_1}(6,5)$ and $V_{i,T_1}(6,4)$ are the values of the (6,5) and (6,4) contracts at the end of year 1 in state i. If the borrower does not prepay in year 1, he obtains a security with value $V_{i,T_1}(6,5)$. The alternative is paying 20% of the principal and converting to 80% of a contract $V_{i,T_1}(6,4)$. Therefore at the end of year 1, the original mortgage has value

$$V_{i,T_1}(7,5) = \min\left(V_{i,T_1}(6,5), \frac{1}{5} + \frac{4}{5}V_{i,T_1}(6,4)\right), \qquad i = 0,\ldots,T_1. \tag{5.3}$$

Knowing the terminal values, discounting back to $t = 0$, taking into account early exercise, is done in the recursion

$$\begin{aligned} V_{it}(7,5) &= \min\left(V'_{it}(7,5), \frac{1}{5} + \frac{4}{5}V_{it}(6,4)\right), \qquad (5.4)\\ &\text{with}\\ V_{it}(6,4) &= d_{it}\left(y + \frac{1}{2}V_{i,t+1}(6,4) + \frac{1}{2}V_{i+1,t+1}(6,4)\right)\\ V'_{it}(7,5) &= d_{it}\left(y + \frac{1}{2}V_{i,t+1}(7,5) + \frac{1}{2}V_{i+1,t+1}(7,5)\right), \end{aligned}$$

where y is the periodic coupon. The prepayment condition compares the continuation value $V'_{it}(7,5)$ with the conversion value $(1 + 4V_{it}(6,4))/5$, and sets $V_{it}(7,5)$ to the minimum of the two.

The problem with this recursion is that the values of the (6,5) and (6,4) contracts are still unknown. We need to expand the annual lattice until we reach nodes with contracts that we can value without path dependence. If, starting at node (6,5), the borrower again does not prepay in year 2, we arrive at node (5,5). This is a contract for which equation 5.2 provides the value. We can compute that value in the lattice by starting at the end of year 7, and discounting back to the end of year 2 each of the five constituent callable bonds. An efficient way to organize all computations will be discussed below.

Nodes in the prepayment conversion lattice recombine: node (5,4) is reached by prepaying in year 1 and not in year 2 and also by prepaying in year 2 and not in year 1. Terminal nodes can be reached in two ways. First, whenever we are at a node (m, m) we have a contract without path dependence which can be priced using equation 5.2. Second, if we reach a node $(m, 1)$ we have a fully callable bond, which can be called in one step. Such a contract can also be easily priced in a spot rate lattice. Formally, the nodes of the prepayment lattice have transitions

$$
(m,n) \begin{array}{ll} \nearrow \ (m-1,n-1) & \text{if } n>1 \\ \\ \searrow \ (m-1,n) & \text{if } n<m. \end{array} \tag{5.5}
$$

Once we have obtained values for all the path independent contracts, we can work backwards to price the original mortgage. The procedure is identical to equation 5.4 above.

For the intermediate periods in a particular calendar year specified by the periods $t \in \left((M-m)\frac{T_M}{M},(M-m+1)\frac{T_M}{M}\right)$, our aim is to find $V_{it}(m,n)$ for all attainable n, given the prepayment conversion lattice as in figure 5.1. Using general notation, the recursion, including partial prepayment decision, is given by

$$
\begin{aligned}
V_{it}(m,n) &= \min\left(V_{it}'(m,n), \frac{1}{n}+\frac{n-1}{n}V_{it}(m-1,n-1)\right) \qquad (5.6)\\
&\text{with}\\
V_{it}'(m,n) &= d_{it}\left(y+\frac{1}{2}V_{i,t+1}(m,n)+\frac{1}{2}V_{i+1,t+1}(m,n)\right)\\
V_{it}(m-1,n-1) &= d_{it}\left(y+\frac{1}{2}V_{i,t+1}(m-1,n-1)+\frac{1}{2}V_{i+1,t+1}(m-1,n-1)\right).
\end{aligned}
$$

For each period a prepayment decision must be made. Before arrival in such a period the contract is characterized by m calendar years and n prepayments to go. No prepayment means continuation of the same contract (m,n). The value is obtained by discounting the expected contract value plus interest payment. Prepayment implies that the remaining contract has $m-1$ calendar years and $n-1$ opportunities to prepay, with value $V_{it}(m-1,n-1)$. The fraction $\frac{1}{n}$ of the contract is repaid (at the nominal price of 1) and the remaining fraction $\frac{n-1}{n}$ is obtained in the contract $(m-1,n-1)$. The value of this contract is also determined by backward recursion. The mortgage value $V_{it}(m,n)$ is the minimum of the continuation value and the portfolio of prepayment value and remaining contract value.

The end-of-year early exercise conditions are very different from the early prepayment boundary of a mortgage with full prepayment. At the end of the year a node (m,n) in the prepayment conversion lattice leads to either $(m-1,n)$ or $(m-1,n-1)$, whereas in other periods (m,n) either remains (m,n) or changes to $(m-1,n-1)$. The recursion for end-of-year nodes (for which $t = (M - m + 1)\frac{T_M}{M}$) is given by

$$\begin{aligned} V_{it}(m,n) &= \min\left(V_{it}(m-1,n), \frac{1}{n} + \frac{n-1}{n}V_{it}(m-1,n-1)\right) \qquad (5.7)\\ &\text{with}\\ V_{it}(m-1,n) &= d_{it}\left(y + \frac{1}{2}V_{i,t+1}(m-1,n) + \frac{1}{2}V_{i+1,t+1}(m-1,n)\right)\\ V_{it}(m-1,n-1) &= d_{it}\left(y + \frac{1}{2}V_{i,t+1}(m-1,n-1) + \frac{1}{2}V_{i+1,t+1}(m-1,n-1)\right). \end{aligned}$$

In any end-of-year node, prepayment takes place if $\frac{1}{n} + \frac{n-1}{n}V_{i,T_l}(m-1,n-1)$ is smaller than the continuation value $V_{i,T_l}(m-1,n)$. This latter value depends on the bonds still to be called with values $C(\ell)$. Prepayment is likely if the continuation value is high, which happens for high $C(\ell)$, so at low interest rates. Low interest rates might give rise to 'December' prepayments to decrease the future unpaid balance, and hence interest payments. A partial prepayment can be optimal even if an immediate loss is faced, as long as the smaller future interest payments offset this loss. 'December' prepayments might be optimal, because a prepayment option expires if not exercised in December. We stress the important result that a partial prepayment can be optimal in nodes where full prepayment is not.

Calendar year ends

Until now we have assumed that calendar years end between two layers of nodes. In that case each node belongs to a single calendar year. From this definition it is clear to which year each prepayment decision belongs. The left panel of figure 5.2 shows calendar years ending between nodes.

Calendar year ends might also fall exactly at interest rate nodes, which is represented by the middle panel of figure 5.2. In these nodes two partial prepayments are allowed, one for the previous year (at December 31 of year ℓ) and one for the coming year (at January 1 of year $\ell + 1$). As a result complexity increases in two ways. First, two prepayment opportunities must be considered at the terminal nodes of each calendar year. Second, even for annual periods a choice must be made in which period to prepay.

One possibility to cope with nodes being part of two successive calendar years is to consider both prepayment opportunities separately. In period $t = (M - m + 1)\frac{T_M}{M}$ we aim

FIGURE 5.2: Calendar year ends

The figure shows calendar year ends in a small lattice example. In the left panel calendar years end between nodes. In the middle panel calendar years end at nodes. The problem of having two prepayment opportunities in these nodes can be circumvented by splitting the nodes and assigning one prepayment opportunity to both adjacent calendar years, as is shown in the right panel.

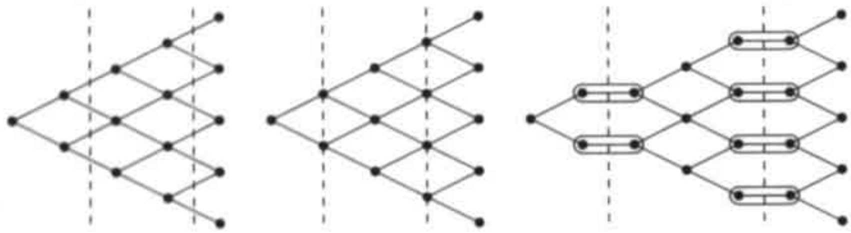

to find $V_{it}(m,n)$ for all attainable n and for all states i. The states in the prepayment conversion lattice reachable from (m,n) are $(m-1,n)$ and $(m-1,n-1)$. For year $\ell+1$, in case of continuation we have $m-1$ calendar years for n prepayments, whereas after prepayment we have $m-1$ calendar years for $n-1$ prepayments. Of course the transitions only hold for $m>n$ and $n>1$, otherwise terminal nodes (m,m) or $(m,1)$ are reached, for which mortgage values are known. For both possibilities we must then consider prepayment in year ℓ. The recursion for these end-of-year nodes, including two prepayment opportunities, is given by

$$\begin{aligned}
V_{it}(m,n) &= \min\left(V_{it}(m-1,n), \frac{1}{n}+\frac{n-1}{n}V_{it}(m-1,n-1)\right) && (5.8)\\
&\text{with} \\
V_{it}(m-1,n) &= \min\left(V'_{it}(m-1,n), \frac{1}{n}+\frac{n-1}{n}V_{it}(m-2,n-1)\right) \\
V_{it}(m-1,n-1) &= \min\left(V'_{it}(m-1,n-1), \frac{1}{n-1}+\frac{n-2}{n-1}V_{it}(m-2,n-2)\right) \\
&\text{and} \\
V'_{it}(m-1,n) &= d_{it}\left(y+\frac{1}{2}V_{i,t+1}(m-1,n)+\frac{1}{2}V_{i+1,t+1}(m-1,n)\right) \\
V_{it}(m-2,n-1) &= d_{it}\left(y+\frac{1}{2}V_{i,t+1}(m-2,n-1)+\frac{1}{2}V_{i+1,t+1}(m-2,n-1)\right) \\
V'_{it}(m-1,n-1) &= d_{it}\left(y+\frac{1}{2}V_{i,t+1}(m-1,n-1)+\frac{1}{2}V_{i+1,t+1}(m-1,n-1)\right)
\end{aligned}$$

$$V_{it}(m-2,n-2) \quad = \quad d_{it}\left(y + \frac{1}{2}V_{i,t+1}(m-2,n-2) + \frac{1}{2}V_{i+1,t+1}(m-2,n-2)\right).$$

Now, the values $V_{it}(m-1,n)$ and $V_{it}(m-1,n-1)$ are not determined only by backward recursion. A second prepayment is allowed, belonging to the start of the next calendar year. If the contract $(m-1,n)$ is partially prepaid, the new contract is $(m-2,n-1)$ as prepayment is not allowed in the remaining of the next calendar year, and one prepayment opportunity is given up. If not, the continuation value is $V_{it}(m-1,n)$ and prepayment is still allowed in the remaining of the next year. A similar argument holds for the contract $(m-1,n-1)$. Note that in case of a second prepayment, the prepayment fraction equals $\frac{1}{n-1}$ and the fraction in the new contract $(m-2,n-2)$ equals $\frac{n-2}{n-1}$.

If calendar year ends occur exactly at nodes, two prepayment decisions must be taken in these nodes. If valuation problems with the calendar year split through the nodes can be solved, problems with calendar years ending between nodes, crossing edges of the lattice, can be solved as well because the prepayment conditions prove to be easier. In the latter case only one prepayment option per time period has to be considered.

In the sequel we will assume calender year ends between layers. This assumption is not restrictive because each tree can be redesigned such that only one prepayment is allowed in each node, even when a node is part of two successive years, as is shown in the right panel of figure 5.2.

5.3 Results

Mortgages with different fixed rate periods and prepayment options are compared on a fair rate basis. The fair rate of a mortgage is the contract rate at which the present value of all payments equals the nominal value. If a mortgage rate is fair, neither the bank nor the client can make a profit, based on the currently observed term structure of interest rates and volatilities.

Fair rates for interest-only mortgage contracts with a varying number of prepayment opportunities N and a varying fixed rate period of M years are presented in tables 5.1 to 5.3. All mortgages have a lifetime of 30 years. Tables differ with respect to the number of periods in the underlying interest rate lattice. We consider monthly, semi-monthly and weekly periods respectively. N takes on the values 1 (corresponding to a fully callable

TABLE 5.1: Fair rates for monthly periods

The table reports fair rates (in annualized percentages) for interest-only mortgages with $M = 5$ or $M = 10$ and $N = 1$, $N = 5$, $N = 10$ or $N \to \infty$. All mortgages mature after 30 years. Underlying term structures on four different dates are considered, with short rates transformed to monthly rates. Depending on M the interest rate lattice requires 60 and 120 periods respectively. The interest payments are based on an annual contract rate following $y = Y/K$ with $K = 12$. The case $N \to \infty$ corresponds to a non-callable mortgage, a mortgage is fully callable if $N = 1$.

	$M = 10$				$M = 5$			
Date N:	∞	10	5	1	∞	10	5	1
Feb 29, 2000	5.856	6.026	6.183	6.568	5.323	5.400	5.504	5.869
Feb 15, 2001	5.305	5.509	5.698	6.143	4.955	5.042	5.167	5.612
June 1, 2001	5.423	5.572	5.715	6.134	4.923	4.996	5.098	5.544
July 2, 2001	5.451	5.609	5.758	6.189	4.874	4.949	5.054	5.504

TABLE 5.2: Fair rates for semi-monthly periods

The table reports fair rates (in annualized percentages) for interest-only mortgages with $M = 5$ or $M = 10$ and $N = 1$, $N = 5$, $N = 10$ or $N \to \infty$. All mortgages mature after 30 years. Underlying term structures on four different dates are considered, with short rates transformed to semi-monthly rates. Depending on M the interest rate lattice requires 120 and 240 periods respectively. The interest payments are based on an annual contract rate following $y = Y/K$ with $K = 24$. The case $N \to \infty$ corresponds to a non-callable mortgage, a mortgage is fully callable if $N = 1$. Full prepayment is allowed at the end of each month.

	$M = 10$				$M = 5$			
Date N:	∞	10	5	1	∞	10	5	1
Feb 29, 2000	5.849	6.021	6.185	6.577	5.317	5.396	5.503	5.892
Feb 15, 2001	5.299	5.503	5.698	6.156	4.950	5.039	5.168	5.664
June 1, 2001	5.417	5.565	5.711	6.140	4.918	4.992	5.096	5.586
July 2, 2001	5.445	5.599	5.750	6.192	4.869	4.946	5.053	5.556

TABLE 5.3: Fair rates for weekly periods

The table reports fair rates (in annualized percentages) for interest-only mortgages with $M = 5$ or $M = 10$ and $N = 1$, $N = 5$, $N = 10$ or $N \to \infty$. All mortgages mature after 30 years. Underlying term structures on four different dates are considered, with short rates transformed to weekly rates. Depending on M the interest rate lattice requires 240 and 480 periods respectively. The interest payments are based on an annual contract rate following $y = Y/K$ with $K = 48$. The case $N \to \infty$ corresponds to a non-callable mortgage, a mortgage is fully callable if $N = 1$. Full prepayment is allowed at the end of each month.

	$M = 10$				$M = 5$			
Date N:	∞	10	5	1	∞	10	5	1
Feb 29, 2000	5.845	6.015	6.179	6.560	5.314	5.393	5.500	5.877
Feb 15, 2001	5.296	5.500	5.696	6.154	4.948	5.035	5.163	5.652
June 1, 2001	5.414	5.566	5.715	6.144	4.915	4.989	5.092	5.572
July 2, 2001	5.442	5.594	5.746	6.180	4.867	4.944	5.051	5.536

mortgage), 5, 10 or infinity (non-callable mortgage). The fixed rate period M equals 5 or 10 years. Calendar year ends fall between nodes, such that each node belongs to a single calendar year. In case a calendar year end falls at a node (implying that two prepayments are allowed in end-of-year nodes, because these nodes are part of two calendar years), only the fair rates for contracts with $M = 10$ and $N = 5$ are affected, but these only increase (prepayment is less restricted) by 0.6 basis point at most.

Increasing the number of layers by decreasing the period length has only a minor effect on the fair rates. Fully callable mortgages exhibit a fair rate difference that can increase to 5 basis points for five year fixed rate periods. Non-callable and partially callable mortgages show differences not larger than 1 basis point when increasing the number of periods.

Tables 5.4 to 5.6 show that a partial prepayment option has significant value, although a prepayment of only a small loan amount is allowed once per calendar year. The fair rate spread of a mortgage is defined as the difference between the fair rate of the mortgage itself and the fair rate of a non-callable mortgage with otherwise the same conditions. For an interest-only mortgage with a ten year fixed rate period, the average full prepayment spread (in terms of fair rates) equals 0.75 percentage point and hardly changes for a finer grid. Hence, the average full prepayment option is worth 75 basis points.

A partial prepayment option, embedded in an interest-only mortgage with a ten year

TABLE 5.4: Fair rate spreads for monthly periods

The table reports fair rate spreads (in annualized percentages) for interest-only mortgages with $M = 5$ or $M = 10$ and $N = 1$, $N = 5$ or $N = 10$, using the non-callable mortgage as benchmark. Underlying term structures on four different dates are considered, with short rates transformed to monthly rates. Depending on M the interest rate lattice requires 60 and 120 periods respectively. The interest payments are based on an annual contract rate following $y = Y/K$ with $K = 12$.

		$M = 10$			$M = 5$		
Date	N:	10	5	1	10	5	1
Feb 29, 2000		0.170	0.327	0.712	0.077	0.181	0.546
Feb 15, 2001		0.204	0.393	0.838	0.087	0.212	0.657
June 1, 2001		0.149	0.292	0.711	0.073	0.175	0.589
July 2, 2001		0.158	0.307	0.738	0.075	0.180	0.630
average		0.170	0.330	0.750	0.078	0.187	0.606

TABLE 5.5: Fair rate spreads for semi-monthly periods

The table reports fair rate spreads (in annualized percentages) for interest-only mortgages with $M = 5$ or $M = 10$ and $N = 1$, $N = 5$ or $N = 10$, using the non-callable mortgage as benchmark. Underlying term structures on four different dates are considered, with short rates transformed to semi-monthly rates. Depending on M the interest rate lattice requires 120 and 240 periods respectively. The interest payments are based on an annual contract rate following $y = Y/K$ with $K = 24$.

		$M = 10$			$M = 5$		
Date	N:	10	5	1	10	5	1
Feb 29, 2000		0.172	0.336	0.728	0.079	0.186	0.575
Feb 15, 2001		0.204	0.399	0.857	0.089	0.218	0.714
June 1, 2001		0.148	0.294	0.723	0.074	0.178	0.668
July 2, 2001		0.154	0.305	0.747	0.077	0.184	0.687
average		0.169	0.334	0.764	0.080	0.192	0.661

TABLE 5.6: Fair rate spreads for weekly periods

The table reports fair rate spreads (in annualized percentages) for interest-only mortgages with $M = 5$ or $M = 10$ and $N = 1$, $N = 5$ or $N = 10$, using the non-callable mortgage as benchmark. Underlying term structures on four different dates are considered, with short rates transformed to weekly rates. Depending on M the interest rate lattice requires 240 and 480 periods respectively. The interest payments are based on an annual contract rate following $y = Y/K$ with $K = 48$.

		$M = 10$			$M = 5$		
Date	N:	10	5	1	10	5	1
Feb 29, 2000		0.170	0.334	0.715	0.079	0.186	0.563
Feb 15, 2001		0.204	0.400	0.858	0.087	0.215	0.704
June 1, 2001		0.152	0.301	0.730	0.074	0.177	0.657
July 2, 2001		0.152	0.304	0.738	0.077	0.184	0.669
average		0.170	0.335	0.760	0.079	0.191	0.648

fixed rate period and $N = 5$, implying a 20% penalty free prepayment each calendar year, has an average premium of 33 basis points. Therefore, a 20% prepayment option is worth 44% of a full prepayment option. The premium for a 10% prepayment option ($N = 10$) is half the premium of a 20% prepayment option: on average 17 basis points or 22% of a full prepayment option.

For shorter fixed rate periods, the effects of prepayment are smaller. When considering a fixed rate period of 5 years, a full prepayment option is worth between 61 and 66 basis points. The premium of a 20% (10%) prepayment option is about 30% (12%) of the premium of a full prepayment option, corresponding to 19 (8) basis points.

Figure 5.3 shows the effect of changing the number of prepayment opportunities on the mortgage price. The underlying interest rate lattice has semi-monthly periods and is calibrated at February 29, 2000. The fixed rate period equals 10 years ($M = 10$). The contract rate is either 6% or 6.5%. For a large contract rate the mortgage value is more sensitive to prepayment options, as the price difference between a fully callable and a non-callable mortgage is larger. Prices increase faster for increasing prepayment opportunities if N is low. For N approaching infinity, the price asymptotically approximates the non-callable mortgage value. The mortgage principal has been scaled to 1.

FIGURE 5.3: Mortgage price for varying N

This figure displays mortgage prices for a varying number of prepayment opportunities N. The February 29, 2000 instance is used with semi-monthly periods and $M = 10$. Results are based on two different contract rates. Both curves asymptotically approach the non-callable mortgage value ($N \to \infty$), represented by the horizontal dashed lines. The nominal loan value equals 1.

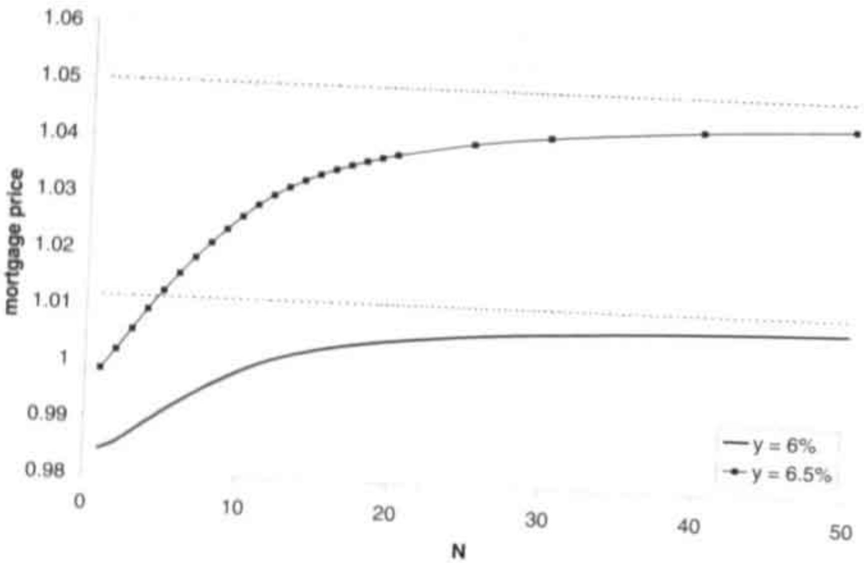

5.4 Concluding remarks

The bond portfolio algorithm to value mortgage contracts with partial prepayments, as described in this chapter, is an efficient algorithm for the valuation of interest-only mortgages, based on a recombining tree approach. The algorithm decomposes an interest-only mortgage contract into several callable bonds and, if necessary, one non-callable bond. However, the applicability of the decomposition method is limited to specific, although popular, contracts.

Prepayment of a partially callable mortgage may occur earlier than a full prepayment, because at the end of a calendar year a prepayment option expires. A partial prepayment in December, just before option expiration, can be optimal, even if a full prepayment is not, thereby decreasing the remaining loan and future interest payments.

The main empirical result of this chapter comprises the importance of a partial prepayment option. In terms of fair rates, a 20% prepayment option is worth 33 basis points on average, when considering a 10 year fixed rate period. This is almost half the value of an unrestricted prepayment option. Including 1% commission costs, a partial prepayment option even proves to be worth 60% of a full prepayment option. Chapter 7 will provide an extensive overview and comparison of fair rates of all types of mortgages and embedded options. The effect of commission costs on the value of prepayment options will be analyzed as well.

An efficient valuation of interest-only mortgages is possible when the number of prepayments N is integer. The majority of Dutch mortgage contracts includes a prepayment option of 10% or 20% of the principal. Mortgages including such option can be valued by the approach proposed in this chapter. However, a small proportion of mortgage contracts includes a 15% prepayment option. The valuation procedure discussed in this chapter can be used to find accurate bounds on the fair rate by solving for $N = 6$ and $N = 7$. Moreover, the bond portfolio algorithm can still be used if the number of prepayment opportunities (i.e. calendar years M) is less than or equal to $\lfloor N \rfloor$. Then the mortgage loan can never be fully amortized in any of the final nodes and all prepayment amounts are maximal.

A complicating factor for the valuation of general mortgages with partial prepayments is the introduction of regular amortization during the contract lifetime. Periodical payments depend on the remaining loan, whereas the prepayment amounts are independent of the loan value. Because not only the number of prepayments determines the remaining loan value, but also the prepayment and redemption amounts, a decomposition into callable bonds is not possible. Redemption amounts are uncertain at the moment of decomposition and depend on the timing of prepayments. Consequently, due to a regular amortization schedule at least one prepayment will be less than the maximally allowed amount. Both the timing and size of this prepayment are uncertain and depend on the underlying term structure. In the next chapter a linear programming model is formulated, which can be used to price general mortgage types and prepayment options. The approach in this chapter is efficient but specific, in the next chapter a general and therefore less efficient approach will be developed.

Chapter 6

Mortgage Valuation with Partial Prepayments[1]

6.1 Introduction

Fully callable mortgages can be priced by a lattice or recombining tree. Mortgages with partial prepayment options are more difficult to price due to path dependencies. Both past and future prepayment decisions affect the current prepayment decision and the current mortgage value. Valuation of partially callable mortgages without regular periodical amortization has been analyzed in the previous chapter, based on a lattice method. For the valuation of partially callable mortgages including regular amortization we must rely on non-recombining tree methods, because lattices are not able to capture path dependencies.[2] Nielsen and Poulsen [70], for example, apply a combination of a lattice and a tree approach to price mortgage contracts with delivery options, introducing path dependencies, where optimality decisions are only taken in a small subset of time periods. Between decision dates the unique scenario path for both mortgage price and interest rate is given. At decision dates the state space behaves as a tree, branching out and not recombining. The number of states increases exponentially with the number of decision dates.

Although the popularity of interest-only and savings mortgages is increasing, traditional mortgage loans, mainly annuities, still make up a significant part of the Dutch mortgage pool. The previous chapter covered the valuation of interest-only mortgages with partial

[1]This chapter is based on joint work with A. Kolen.

[2]Another alternative, Monte Carlo simulation has difficulties to value American option types.

prepayments. Now we will focus on the valuation of annuity mortgages. Each month a part of the loan is redeemed to the mortgagee. An annuity has constant periodic payments, with increasing redemption amounts and decreasing interest payments.

Dutch mortgage contracts allow a prepayment of 10% to 20% of the principal per calendar year. Due to the regular periodic payments, optimal prepayment behavior differs from prepayment of interest-only mortgages. After a prepayment the amortization schedule changes, which can be a reason to postpone this prepayment. Part of an optimal prepayment strategy can be to reduce the future unpaid balance. A later prepayment reduces this unpaid balance more than an early prepayment, as the regular amortization schedule reduces the unpaid balance more before an additional prepayment. This creates an incentive to postpone prepayment relative to a mortgage without prepayment restrictions. Interest-only mortgages are not subject to this effect, as the unpaid balance of these contracts does not decrease by means of a regular amortization schedule.

Another difference between the valuation of fully and partially callable mortgages is the existence of 'December' prepayments for partial prepayment options. We have concluded in the previous chapter that a partial prepayment might occur earlier than a full prepayment since otherwise a prepayment option expires with the ending of a calendar year. This effect holds for all partially callable mortgages, independent of the amortization schedule.

In this chapter we formulate a linear programming model for the valuation of partially callable annuity mortgages. The LP formulation can also capture linear and interest-only mortgages. All time periods in our model allow for prepayment of a part of the mortgage loan, involving the use of a complete non-recombining tree. Linear programming is applied both to obtain an exact mortgage value and prepayment strategy and, using duality theory, to derive bounds on the optimal mortgage value. The next section introduces the mathematical framework, section 6.3 builds the LP model. The dual problem is formulated in section 6.4. The implications of the LP formulation for fully callable mortgages are provided in section 6.5, based on duality theory. Section 6.6 solves an accurate heuristic for the original LP model, obtaining upper bounds on the mortgage price and lower bounds on the fair rate. We also narrow the gap between upper and lower bound on the mortgage price, in order to derive an accurate approximation. Results are provided in section 6.7.

6.2 Mathematical framework

The prepayment strategy of any mortgage with partial prepayment opportunities is path dependent because the optimality of a current prepayment decision depends on past and future prepayment decisions. Periodical payments depend on the adopted strategy. Moreover, the allowance of a prepayment decision depends on whether a prepayment occurred in the same calendar year. The prepayment decision itself depends on the profitability of postponing a prepayment to a later month in the same calendar year.

The problem is formulated on a non-recombining tree. The states in a non-recombining tree are labelled as in figure 6.1. The root node has label 0, the two nodes at time 1 are labelled 1 and 2. Generally, the transitions are given by

$$i \begin{array}{l} \nearrow\ 2i+1 \\ \\ \searrow\ 2i+2. \end{array} \tag{6.1}$$

The time period $t(i)$ corresponding to state i is

$$t(i) = \lfloor {}^2\log(i+1) \rfloor. \tag{6.2}$$

The final period $t(i) = T$. The final node is labelled $m = 2^{T+1} - 2$. Note that the unique predecessor of state i, if not the root node, is labelled $\lfloor (i-1)/2 \rfloor$. A state i, for which $t(i) = T$, is called a leaf node. All nodes that are neither the root node nor a leaf node are called intermediate nodes.

We will focus on mortgage contracts with a regular redemption schedule, of which an annuity mortgage is most popular and well-known. An annuity contract is characterized by a principal amount P, a periodic contract rate y and a maturity L. A constant regular payment occurs at all $t(i) = 1, \ldots, L$. The size of this payment M_i in state i depends on the unpaid balance in the previous period, which in turn depends on the past prepayment behavior, and on the contract rate and remaining lifetime $n_i = L - t(i)$. The unpaid balance U_i in state i is defined as the amount of money still owed to the mortgagee at this state. By definition $U_0 = P$. The periodic payment is now given by

$$\begin{aligned} M_i &= U_{\lfloor (i-1)/2 \rfloor} \cdot f(y, n_i + 1) \\ &\text{with} \\ f(y, n_i) &= \frac{y}{1-(1+y)^{-n_i}}. \end{aligned}$$

FIGURE 6.1: Non-recombining binomial tree

The figure shows the first four time steps of a non-recombining binomial tree and the applied node labelling. Transitions have probability $\frac{1}{2}$.

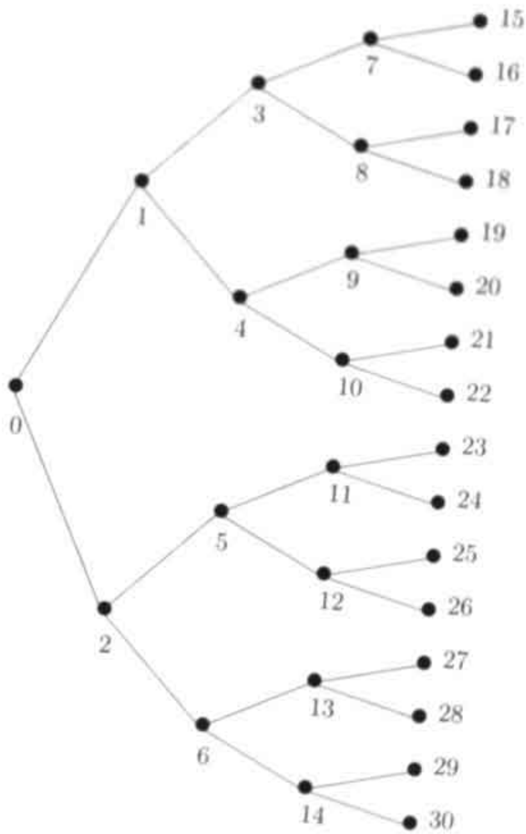

In case of no prepayments other than the periodic payments, the periodic payment is constant over all states, that is $M_i = M$. Linear and interest-only mortgages are easily modelled by only adapting $f(y, n_i)$ according to the mortgage specifications in chapter 4.

An annuity mortgage with partial prepayments is defined by the principal amount P, a periodic contract rate y, a maturity L and K consecutive subintervals I_k of $[t_0, \ldots, t_n]$, containing all distinct $t(i)$ in increasing order. The endpoints of all subintervals belong to the set $\{t_0, \ldots, t_n\}$ such that each point in the set is in exactly one interval. The first set starts at t_0, the last set ends with t_n. We consider a single fixed rate period, ending at t_n. In each interval I_k the total amount that can be repaid in this interval is restricted to be less than or equal to F_k. In most cases all intervals have equal length (for example a calendar year) and the prepayment is restricted to at most a fixed percentage α of the principal amount: $X \equiv F_k = \alpha \cdot P$ for all k.

The actual prepayment in state i is denoted by x_i. At the end of each fixed rate period, the remaining loan balance can be fully repaid without penalty. The mortgage price in leaf nodes is therefore equal to the remaining unpaid balance. Consequently, in the optimization model the prepayment amount x_i can be set to zero in all leaf nodes. In fact, in leaf nodes the mortgagor is indifferent to prepay.

The interest rate process is given by a one-period interest rate on all nodes in the state space, denoted by r_i. An interest rate scenario is represented by a path from the root node to a leaf. A state price λ_i is the root price of a security that pays out 1 in state i and zero in all other states. The state price is recursively defined by

$$\lambda_i = \frac{\lambda_{\lfloor (i-1)/2 \rfloor}}{2(1 + r_{\lfloor (i-1)/2 \rfloor})}, \quad i = 1, \ldots, m$$

and $\lambda_0 = 1$. State prices are used for discounting cash flows along a scenario path. The present value of an asset, paying a cash flow of c_i in state i and zero in all other states, is equal to $\lambda_i c_i$.

6.3 The model

The model formulation is based on Kolen [53]. We assume that no prepayment takes place at $t = 0$ (this could be accounted for in the initial loan amount). Formally,

$$U_0 = P. \tag{6.3}$$

Every period the unpaid balance increases at rate y. A regular amount M_i, including interest and redemption amounts, is paid in state i. Additionally, the mortgagor must decide whether to prepay an amount up to the allowed $X = \alpha \cdot U_0$, with α the maximally allowed prepayment percentage. The first class of constraints models the unpaid balance in an intermediate state i:

$$U_i = U_{\lfloor (i-1)/2 \rfloor} \cdot (1+y) - U_{\lfloor (i-1)/2 \rfloor} \cdot f(y, n_i + 1) - x_i, \quad i = 1, \ldots, m/2 - 1. \tag{6.4}$$

Because no prepayment occurs in a final state, the unpaid balance in such state equals

$$U_i = U_{\lfloor (i-1)/2 \rfloor} \cdot (1+y) - U_{\lfloor (i-1)/2 \rfloor} \cdot f(y, n_i + 1), \quad i = m/2, \ldots, m.$$

The next class of constraints models the upper bound on the total prepayments within a given time interval. Let us denote by Q_k the set of all paths for which the first state on the path belongs to the layer corresponding to the begin point of interval I_k and the last state on the path belongs to the layer corresponding to the end point of the interval I_k, $k = 1, \ldots, K$. Then the additional prepayment amount x_i is restricted by

$$\sum_{i \in Q} x_i \leq X, \qquad Q \in Q_k, \qquad k = 1, \ldots, K. \tag{6.5}$$

We consider a constant prepayment amount $X = F_k$ and all subintervals make up exactly one calendar year.

An optimal prepayment strategy for the mortgagor is the strategy that minimizes the present value of all payments. These payments include regular payments M_i, additional payments x_i and, if any, redemption of the remaining contract value at the leaf nodes. These latter values equal the unpaid balances at the penultimate states, increased by the contract interest rate (recall that no prepayment occurs at leaf nodes). Payments are discounted by means of the state prices λ_i. The mortgage value is now represented by

$$V_0 = \sum_{i=1}^{m/2-1} \lambda_i [U_{\lfloor (i-1)/2 \rfloor} \cdot f(y, n_i + 1) + x_i] + \sum_{i=m/2}^{m} \lambda_i U_{\lfloor (i-1)/2 \rfloor}(1+y). \tag{6.6}$$

Now, the linear programming objective for pricing annuity mortgages with partial prepayments is to

$$\text{minimize} \qquad \sum_{i=1}^{(m-2)/2} \lambda_i [U_{\lfloor (i-1)/2 \rfloor} \cdot f(y, n_i + 1) + x_i] + \sum_{i=m/2}^{m} \lambda_i U_{\lfloor (i-1)/2 \rfloor}(1+y)$$

subject to

$$\begin{aligned}
U_0 &= P \\
U_i &= U_{\lfloor (i-1)/2 \rfloor} \cdot (1+y) - U_{\lfloor (i-1)/2 \rfloor} \cdot f(y, n_i + 1) - x_i, \quad i = 1, \ldots, (m-2)/2 \\
\sum_{i \in Q} x_i &\leq X, \quad \forall Q \in Q_k, \quad \forall k \\
U_i &\geq 0, \quad \forall i \\
x_i &\geq 0, \quad \forall i
\end{aligned}$$

The last two restrictions state that the borrower can never prepay more than the unpaid balance and that prepaid amounts cannot be taken out again.[3]

An upper bound on the mortgage value can be obtained by constructing a feasible solution to the general (primal) problem. No prepayment, equating all x variables to zero, is a trivial feasible solution for which the objective boils down to discounting future regular periodical payments. Consequently, the value of a non-callable mortgage is a trivial upper bound on the value of a partial prepayment mortgage with the same contract rate and time to maturity. In order to find the mortgage value with partial prepayments, the variables x_i (and the resulting U_i) of this LP model must be optimized.

As an example, consider a problem instance defined on the state space given in figure 6.1. We assume that we have two time intervals, $I_1 = [t_0, t_1]$ and $I_2 = [t_2, t_3]$.[4] Furthermore, we face a constant maximum prepayment percentage X and a contract lifetime of four periods, that is, the final tree period marks the end of the contract. The model is given below in standard format.[5]

[3]Some mortgage contracts allow taking out earlier prepaid loan amounts. In that case x_i can be restricted to be larger than minus the sum of all previous prepayments, or larger than some contract specification restricting the maximal amount to take back.

[4]Time intervals with the year split through nodes, such that one time period belongs to both the previous and the upcoming year, requires two prepayment variables for each end-of-calendar year node. This can be achieved by assigning one of the prepayments to each of the edges incident to the end-of-calendar year node. For the purpose of the example, this would complicate the formulations and increase the number of variables unnecessarily.

[5]Note that for this small scale example the contract is fully amortized at the end of the fixed rate period (that is, at t_4). In case the mortgage lifetime is longer than the fixed rate period, an analogous formulation can be applied, only changing n_i.

$$\begin{aligned}
\text{minimize} \quad & \lambda_1[U_0 \cdot f(y,4) + x_1] + \ldots + \lambda_{14}[U_6 \cdot f(y,2) + x_{14}] \\
+ \quad & \lambda_{15}U_7(1+y) + \ldots + \lambda_{30}U_{14}(1+y)
\end{aligned}$$

subject to

$$\begin{aligned}
U_0 &= P \\
U_1 - U_0(1+y) + U_0 f(y,4) + x_1 &= 0 \\
&\vdots \\
U_{14} - U_6(1+y) + U_6 f(y,2) + x_{14} &= 0 \\
-x_1 &\geq -X \\
-x_2 &\geq -X \\
-x_3 - x_7 &\geq -X \\
&\vdots \\
-x_6 - x_{14} &\geq -X \\
U_i &\geq 0, \quad \forall i \\
x_i &\geq 0, \quad \forall i
\end{aligned}$$

6.4 Dual formulation

In order to make statements about the optimality of a solution, we apply duality theory. Before deriving the general dual problem formulation, we provide the dual of the example at the end of the previous section. Dual variables v_i and z_i are introduced, the first correspond to the constraint set 6.4, the second to the restrictions 6.5. For each state of the tree there exists one v_i, while a z_i is required only for periods concluding a calendar year as these determine the number of calendar year restrictions 6.5. Denote the set of nodes concluding the calendar years by C. Both v_i and z_i have labels equal to the corresponding state, such that the z_i labels are not continuous. For instance, in the state space example z_3 does not exist because state 3 does not mark a calendar year end. The dual formulation of the four-period example is given by

$$\begin{array}{ll} \text{maximize} & -X\sum_{i\in C} z_i + v_0 P \\ \text{subject to} & \end{array}$$

$$\begin{aligned}
v_0 - v_1(1+y) + v_1 f(y,4) - v_2(1+y) + v_2 f(y,4) &\leq \lambda_1 f(y,4) + \lambda_2 f(y,4) \\
v_1 - v_3(1+y) + v_3 f(y,3) - v_4(1+y) + v_4 f(y,3) &\leq \lambda_3 f(y,3) + \lambda_4 f(y,3) \\
v_2 - v_5(1+y) + v_5 f(y,3) - v_6(1+y) + v_6 f(y,3) &\leq \lambda_5 f(y,3) + \lambda_6 f(y,3) \\
&\vdots \\
v_6 - v_{13}(1+y) + v_{13} f(y,2) - v_{14}(1+y) + v_{14} f(y,2) &\leq \lambda_{13} f(y,2) + \lambda_{14} f(y,2)
\end{aligned}$$

$$\begin{aligned}
v_7 &\leq (\lambda_{15} + \lambda_{16})(1+y) \\
&\vdots \\
v_{14} &\leq (\lambda_{29} + \lambda_{30})(1+y) \\
-z_1 + v_1 &\leq \lambda_1 \\
-z_2 + v_2 &\leq \lambda_2 \\
-z_7 - z_8 + v_3 &\leq \lambda_3 \\
-z_9 - z_{10} + v_4 &\leq \lambda_4 \\
-z_{11} - z_{12} + v_5 &\leq \lambda_5 \\
-z_{13} - z_{14} + v_6 &\leq \lambda_6 \\
-z_7 + v_7 &\leq \lambda_7 \\
&\vdots \\
-z_{14} + v_{14} &\leq \lambda_{14} \\
z_i &\geq 0, \quad \forall i \in C
\end{aligned}$$

Let $C_i \subset C$ denote the set containing all states marking the end of the calendar year to which state i belongs that are attainable from state i. For instance, considering intermediate state 3, $C_3 = \{7, 8\}$. Also, define the function $g(i)$ to be

$$g(i) = (v_{2i+1} + v_{2i+2})(1 + y - f(y, n_i)) + (\lambda_{2i+1} + \lambda_{2i+2}) f(y, n_i).$$

Final period states $i = (m-2)/2, \ldots, m$ have $v_i = \lambda_i$, which can be observed when including the balance constraints $U_i \geq 0$ for these states explicitly in the problem formulation and rewriting the objective to include the remaining unpaid balance U_i for leaf nodes separately, discounted by λ_i. For penultimate states, $g(i)$ can therefore be simplified to

$$g(i) \;=\; (\lambda_{2i+1} + \lambda_{2i+2})(1+y) \;=\; \lambda_i \frac{1+y}{1+r_i}.$$

Now, the complete definition of the function $g(i)$ is

$$g(i) = \begin{cases} (v_{2i+1} + v_{2i+2})(1 + y - f(y, n_i)) + (\lambda_{2i+1} + \lambda_{2i+2}) f(y, n_i), \\ \hfill \forall i = 0, \ldots, (m-6)/4 \\ \\ (\lambda_{2i+1} + \lambda_{2i+2})(1+y), \qquad \forall i = (m-2)/4, \ldots, (m-2)/2. \end{cases} \tag{6.7}$$

The general formulation of the dual problem to value annuity mortgages with partial prepayments is the following:

$$\begin{array}{lrcl} \text{maximize} & -X \sum_{i \in C} z_i + v_0 P & & \\ \text{subject to} & & & \\ & v_i & \leq & g(i), \quad \forall i = 0, \ldots, (m-2)/2 \\ & -\sum_{\ell \in C_i} z_\ell + v_i & \leq & \lambda_i, \quad \forall i = 1, \ldots, (m-2)/2 \\ & z_i & \geq & 0, \quad \forall i \in C \end{array} \tag{6.8}$$

Complementary slackness conditions can be used to find dual variables based on the primal solution. If a primal inequality contains slack, then the corresponding dual variable equals zero. For the restrictions in our mortgage valuation problem, this implies:

$$\sum_{i \in Q} x_i < X \Rightarrow z_\ell = 0, \tag{6.9}$$

where ℓ is the last node, at the time interval end, of path Q. Typically, prepayment is restricted per calendar year, such that path Q covers one year. Node ℓ is then the last node of the year. Condition 6.9 states that if prepayment during scenario path interval Q is less than maximal, the dual variable z_ℓ can be fixed to 0.

When the dual solution is known, complementary slackness can be used to obtain a partial solution to the primal LP problem:

$$z_\ell > 0 \Rightarrow \sum_{i \in Q} x_i = X. \tag{6.10}$$

This complementary slackness condition states that if the dual variable z_ℓ, belonging to state ℓ, is positive, then prepayment is maximal along path Q, which ends in node ℓ and covers exactly one year.

Complementary slackness conditions with respect to the inequalities of the dual formulation can be derived similarly. These conditions, $\forall i = 1, \ldots, (m-2)/2$, read

$$-\sum_{\ell \in C_i} z_\ell + v_i < \lambda_i \quad \Rightarrow \quad x_i = 0 \tag{6.11}$$

$$x_i > 0 \quad \Rightarrow \quad -\sum_{\ell \in C_i} z_\ell + v_i = \lambda_i \tag{6.12}$$

and

$$v_i < g(i) \quad \Rightarrow \quad U_i = 0 \tag{6.13}$$

$$U_i > 0 \quad \Rightarrow \quad v_i = g(i) \tag{6.14}$$

From the dual formulation 6.8 follows that the dual variables v_i must be both less than or equal to $g(i)$ and $\lambda_i + \sum_{\ell \in C_i} z_\ell$. As v_0 (the dual variable to be maximized) is determined by a backward recursion approach depending on all future v_i, we may state that $v_i = \min(g(i), \lambda_i + \sum_{\ell \in C_i} z_\ell), \forall i = 1, \ldots, (m-2)/2$. Hence for given z, the complete dual solution and the corresponding mortgage value can be obtained by backward recursion. The optimal prepayment strategy in state i can be partly derived from this minimum evaluation to obtain v_i, as will be shown by the next two theorems.

Theorem 6.1 *If $\lambda_i + \sum_{\ell \in C_i} z_\ell < g(i)$, then a final prepayment of the remaining loan is optimal in state i.*

Proof Suppose that $\lambda_i + \sum_{\ell \in C_i} z_\ell < g(i)$. Then $v_i = \lambda_i + \sum_{\ell \in C_i} z_\ell < g(i)$, and $U_i = 0$ because of complementary slackness condition 6.13. A full prepayment of the remaining loan is optimal. Similarly, if full prepayment is not optimal in state i, then $U_i > 0$. By complementary slackness condition 6.14, $v_i = g(i)$, which can only be true if $g(i) \leq \lambda_i + \sum_{\ell \in C_i} z_\ell$. ∎

Theorem 6.2 *If $g(i) < \lambda_i + \sum_{\ell \in C_i} z_\ell$, then no positive prepayment of a (partially) callable mortgage is optimal in state i.*

Proof Suppose that $g(i) < \lambda_i + \sum_{\ell \in C_i} z_\ell$. Then $v_i = g(i) < \lambda_i + \sum_{\ell \in C_i} z_\ell$, and $x_i = 0$ because of complementary slackness condition 6.11. No prepayment is optimal. Similarly, if a positive prepayment is optimal in state i, then $x_i > 0$. By complementary slackness condition 6.12, $v_i = \lambda_i + \sum_{\ell \in C_i} z_\ell$, which can only be true if $\lambda_i + \sum_{\ell \in C_i} z_\ell \leq g(i)$. ■

As a direct result from complementary slackness, the theorems imply that for a non-final partial prepayment,

$$\lambda_i + \sum_{\ell \in C_i} z_\ell = g(i)$$

must hold. The theorems on optimal prepayment are difficult to use for partially callable mortgages, since all non-final partial prepayment decisions cannot be determined by either $\lambda_i + \sum_{\ell \in C_i} z_\ell < g(i)$ or $\lambda_i + \sum_{\ell \in C_i} z_\ell > g(i)$. Both theorems are easier applied to fully callable mortgages.

6.5 Implications for fully callable mortgages

Mortgage valuation including full prepayment is a relaxation of the original problem formulated in section 6.3, omitting the limited prepayment restriction. Stated differently, the *maximum* prepayment amount X is infinite for fully callable mortgages. *Actual* prepayments must still satisfy the conditions

$$\begin{aligned} x_i &\geq 0, \quad \forall i \\ U_i &\geq 0, \quad \forall i. \end{aligned}$$

As a result, $\sum_{i \in Q} x_i < X$, $\forall Q$ is a valid constraint for fully callable mortgages as well, assuming X to be infinitely large. By complementary slackness condition 6.9,

$$z_\ell = 0, \ \forall \ell \in C_i, \ \forall i. \tag{6.15}$$

The equations with respect to the dual variables v_i follow directly from the dual programming formulation and the fact that $z_i = 0$, $\forall i \in C$. Therefore, the value of a fully callable mortgage is equal to the dual objective $v_0 P$, where v_0 is given by

$$v_0 = g(0), \tag{6.16}$$

$$v_i = \min\left(g(i), \lambda_i\right), \quad \forall i = 1, \ldots, (m-2)/2. \tag{6.17}$$

Terminal values to the backward recursion of v_i are provided at the penultimate states, at which v_i only depends on state prices and the contract rate, according to the definition of $g(i)$ in 6.7. This approach is comparable to the standard backward recursion (recursively solving the primal LP) applied for the valuation of American options.

Optimal prepayment conditions for a fully callable mortgage are based on complementary slackness and can be easily derived from the theorems on optimal prepayment in the previous section. The optimal prepayment strategy of a fully callable mortgage depends solely on $g(i)$ and the state prices λ_i, according to 6.17. Theorem 6.1 implies that full prepayment of a fully callable mortgage is optimal in state i if $\lambda_i < g(i)$.[6] No positive prepayment of a fully callable mortgage is optimal in state i if $g(i) < \lambda_i$, according to theorem 6.2.

Any dual feasible solution provides a lower bound on the mortgage value. Consequently, the value of a mortgage contract with partial prepayments is bounded from below by the full prepayment value of a mortgage with the same contract rate and time to maturity. The lower bound can be improved by increasing z_i for some i. Although this decreases the lower bound directly, v_i (and by backward recursion 6.17, v_0) can increase due to constraint relaxation. If the increase in v_0P is larger than the rise of $X\sum_{i\in C} z_i$, raising some z_i can improve the dual solution and hence the lower bound on the mortgage price.

Since the problem formulation is based on a non-recombining tree only small problem instances can be solved to optimality. For long term, partially callable mortgage contracts the optimal prepayment strategy cannot be determined efficiently. The next section introduces a heuristic to derive the optimal prepayment strategy based on a lattice approach. This approximative strategy is used to obtain a bound on the mortgage price and on the fair rate.

6.6 Bounding the fair rate

Small problem instances can be solved exactly by either primal or dual formulation, based on a non-recombining tree approach. For large instances (our typical problem size equals 120 periods, resulting in 2^{120} final states), such formulation is not efficiently solvable.

[6]Note that if full prepayment is optimal in state i, the short rate r_i must be smaller than the contract rate y. Hence, we have also shown that if $\lambda_i < g(i)$, then $r_i < y$ (which also follows from the definition of $g(i)$, the recursive defining of the state prices and the restrictions of the dual problem). The converse, if $r_i < y$, then full prepayment is optimal, does not necessarily hold.

Therefore, we focus on obtaining upper and lower bounds on the mortgage value by constructing primal or dual feasible solutions respectively. Any primal feasible solution (that is, an allowed prepayment strategy) implies an upper bound on the optimal value of a partially callable mortgage, or by an iterative procedure, a lower bound on the fair rate. This section constructs a primal feasible solution, based on a lattice approach to retain computational efficiency.

The size of the original lattice equals the length of the first fixed rate period. During this period a large number of prepayment decisions must be taken. Each prepayment originates a new mortgage loan with a smaller unpaid balance, periodical payment and time to maturity. These new mortgage loans are valued by a sublattice of the original lattice, using the corresponding interest rates.

Figure 6.2 shows the decomposition process based on the full prepayment boundary. This boundary is derived according to the optimal prepayment strategy of a fully callable mortgage. Valuation of fully callable mortgages and the derivation of the full prepayment boundary has been discussed in chapter 4. All nodes below the full prepayment boundary are considered as states in which full prepayment (if allowed) is optimal. Full prepayment is not optimal in nodes above the full prepayment boundary. Postponing (a part of) a full prepayment is never optimal. As soon as calling a fully callable mortgage is profitable, the mortgage is fully prepaid. Postponing prepayment leads to higher total interest payments.

Optimal prepayment of a partially callable mortgage can be both earlier and later than an optimal full prepayment. It might be optimal to postpone a partial prepayment if only limited prepayment is allowed. The reason is that higher interest payments are compensated by a lower future unpaid balance, because regular redemption is larger before than after an additional prepayment. A lower unpaid balance leads to lower future periodical payments. If these lower payments (more than) offset the disadvantageous higher interest payments due to postponing prepayment, a later prepayment might be optimal. Consequently, for a partially callable mortgage 'no prepayment' can be the optimal decision in a node below the full prepayment boundary. For fully callable mortgages the unpaid balance after full prepayment is zero, the resulting periodical payments are zero as well, and these payments can therefore not be used as compensation for higher interest payments.

An optimal prepayment strategy might also involve a partial prepayment in a node above the full prepayment boundary. Such an early prepayment can be optimal in December to exercise a prepayment option just before the end of a calendar year, the option expiration date. An extra prepayment reduces the future unpaid balance and periodical

FIGURE 6.2: Decomposition based on full prepayment boundary

The figure shows the main lattice and one of the first level sublattices after a decomposition based on the full prepayment boundary (the horizontal dashed line). All encircled nodes are candidate prepayment nodes. All solid encircled nodes are nodes in which a first prepayment is considered and from which a new sublattice is constructed. Vertical dashed lines represent calendar years. The effective prepayment boundary (longer dashes) is a combination of the full prepayment boundary and one of the calendar year restrictions. Prepayment in the first candidate prepayment node (that is, the root node of the sublattice) implies that the next prepayment cannot be in the same calendar year.

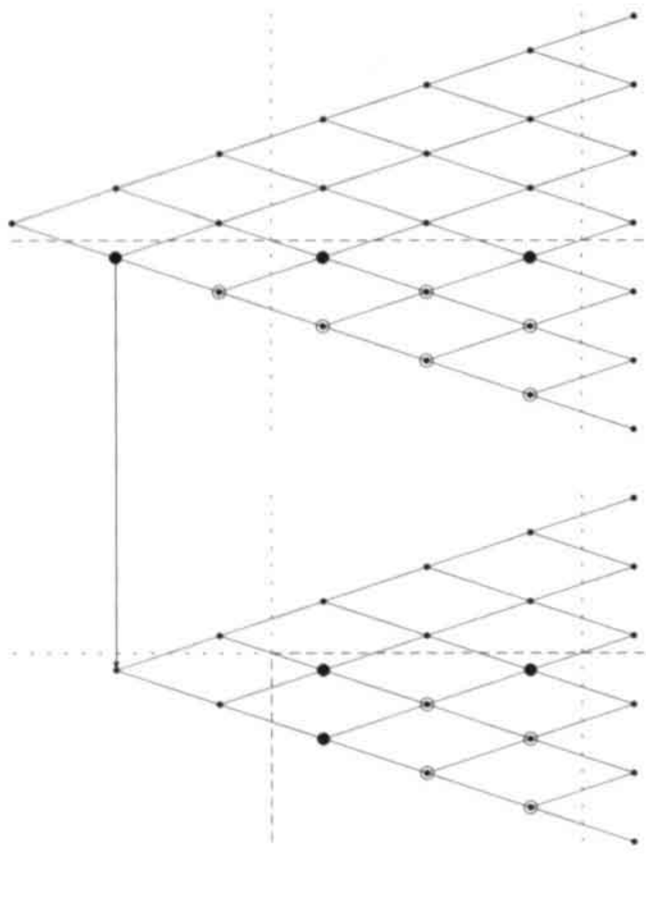

payments. If the resulting lower payments more than offset the disadvantageous prepayment in December, an early partial prepayment can be optimal. 'December' prepayments have been discussed in chapter 5.

Since an optimal partial prepayment can be both earlier and later than an optimal full prepayment, the full prepayment boundary provides a feasible prepayment strategy, but not necessarily the optimal strategy. To construct a primal feasible solution we assume that no prepayment occurs in nodes above the full prepayment boundary and a partial prepayment occurs in nodes below the full prepayment boundary. Additionally, we assume that a partial prepayment amount is always equal to the maximally allowed amount, unless the remaining loan is smaller than the maximal prepayment. In the latter case we assume a final prepayment of the remaining loan.[7]

Our approximation to the optimal prepayment strategy involves no prepayment in nodes above the full prepayment boundary. This part of the valuation process can be performed by a single lattice approach. Furthermore, a maximally allowed prepayment ($x_i = X$) is included whenever the full prepayment boundary is crossed downwards. After each prepayment a new sublattice is constructed based on the remaining mortgage lifetime and unpaid balance. The prepayment boundary in each sublattice is similar to the boundary in the original tree, except for prepayment to start at the first month of a new calendar year. The prepayment node in the parent lattice is the root node of the sublattice.

One of the first level sublattices (after the first partial prepayment), including full prepayment boundaries adapted for calendar year restrictions, is depicted in figure 4.1. The number of levels of sublattices is equal to the maximum number of prepayments. In case of prepayments limited to 20%, the number of levels is bounded by five. The number of sublattices increases with rate T per level. Denote the number of levels by K. A recursion through each sublattice to determine the mortgage price requires a computation time of $O(T^2)$, implying a total computational effort of $O(T^{K+2})$.

Although computation time is of a polynomial order (compared to exponential for a non-recombining tree), the polynomial degree is still large. Efficiency can be improved by performing a recursion only once for all sublattices rooted in the same node. Suppose

[7]In some states the optimal prepayment can be less than maximal. A less than maximal prepayment is optimal if a fully amortizing loan cannot be repaid by periodical payments plus an integer number of additional prepayments. At least one of the additional prepayments is smaller than the maximally allowed amount. This prepayment optimally takes place whenever interest rates are favorable, but less favorable than in the case of a maximally allowed prepayment (which can be in any period, depending on interest rate behavior).

node (i,t) can be reached by two different paths. For the first path a recursion is required to determine the price P_1 corresponding to unpaid balance U_1 in node (i,t). The unpaid balance according to the second path reaching node (i,t) equals U_2. Now the price can be scaled to be $P_2 = U_2 \cdot P_1/U_1$. This pricing strategy is similar to the valuation of an adjustable rate mortgage discussed in section 4.3.2, based on Kau, Keenan, Muller and Epperson [48]. However, additional prepayments are not scalable since these depend on the initial loan and not on the remaining loan. These cash flows are excluded from the traditional valuation procedure, but added separately and discounted at the appropriate discount factors. The scaling approach is more efficient than the standard approach as long as the decrease in the number of recursions is not outweighed by the preprocessing phase of calculating discount factors. This is typically the case for large instances with many prepayment opportunities. Computation time for the scalable decomposition method is of $O(T^4)$, since at most one recursion of $O(T^2)$ is required for each node.[8]

Partially callable mortgages with a fixed rate period of five years can be valued by the (scalable) decomposition method based on the full prepayment boundary, providing an upper bound on the mortgage price or a lower bound on the fair rate. Since many lattices must be stored in memory simultaneously for large instances, loans with ten year fixed rate periods can only be valued when lowering the prepayment boundary significantly. As a result, fewer sublattices have to be constructed. The lower bound on the fair rate will deviate more from the optimal fair rate, as the prepayment strategy differs significantly from the optimal prepayment strategy.[9]

The next section provides lower bounds on the optimal fair rate based on the lattice decomposition method. We are interested in the difference between the optimal fair rate and the calculated lower bound. For this reason, we use the fair rate of an interest-only mortgage (derived in chapter 5) as an upper bound on the fair rate of a partially callable annuity. Since all term structures are upward sloping, an interest-only mortgage faces an unattractive redemption schedule. The fair rate of an interest-only mortgage is therefore higher than the fair rate of an annuity or linear mortgage with similar characteristics.

[8]Fair rates following from the scalable decomposition method slightly differ from fair rates according to the standard decomposition method. Prepayment in node (i,t) according to the standard decomposition method is based on the unpaid balance and price in node (i,t) of the parent lattice. The scalable method, having no recursion in most (sub)lattices and therefore no truly optimal strategy of consecutive prepayments, can only compare unpaid balance and price at the root of the child lattice. The standard decomposition method is more accurate, although fair rate differences are negligible.

[9]We have restricted prepayment opportunities in various ways to improve efficiency. One could choose for allowing prepayment only once or twice per calendar year. However, shifting the prepayment boundary downwards provided the best lower bound on the fair rate.

TABLE 6.1: Fair rates for a five year fixed rate period.

This table provides lower bounds on fair rates of partially callable annuity and linear mortgages. Upper bounds correspond to fair rates of partially callable interest-only mortgages. The underlying interest rate lattice consists of monthly periods. The term structure model is based on a one-factor BDT model. Mortgage contracts have a five year fixed rate period and exclude commission, respectively include a 1% commission.

		no commission		1% commission	
Type	Date	LB	UB	LB	UB
Annuity	Feb 29, 2000	5.48	5.50	5.17	5.21
	Feb 15, 2001	5.14	5.17	4.82	4.85
	June 1, 2001	5.10	5.10	4.77	4.80
	July 2, 2001	5.05	5.05	4.72	4.75
Linear	Feb 29, 2000	5.46	5.50	5.14	5.21
	Feb 15, 2001	5.13	5.17	4.79	4.85
	June 1, 2001	5.08	5.10	4.74	4.80
	July 2, 2001	5.03	5.05	4.69	4.75

6.7 Results

Lower bounds on the fair rates are calculated for both 5 and 10 year fixed rate periods. We consider partially callable mortgages excluding commission and including a 1% commission on four dates. Annuity and linear mortgages are included. A practical upper bound on the fair annuity and linear mortgage rates is obtained by the fair rate of a corresponding interest-only mortgage, since the underlying term structures are upward sloping. The bounds define a range for the optimal fair rate of partially callable annuity and linear mortgages.

For five year fixed rate periods no computational problems arise. When prepaying the maximally allowed amount in any node below the full prepayment boundary and not prepaying anything in any node above, a tight lower bound on the fair rate is obtained. As can be concluded from table 6.1, the lower bound differs between zero and 7 basis points from the upper bound, defined by the fair rate of an interest-only mortgage with similar conditions. Therefore, the lower bound is a very accurate approximation of the optimal fair contract rate. Also, the optimal prepayment strategy will not differ largely from the

TABLE 6.2: Fair rates for a ten year fixed rate period.

This table provides lower bounds on fair rates of partially callable annuity and linear mortgages. Upper bounds correspond to fair rates of partially callable interest-only mortgages. The underlying interest rate lattice consists of monthly periods. The term structure model is based on a one-factor BDT model. Mortgage contracts have a ten year fixed rate period and exclude commission, respectively include a 1% commission.

		no commission		1% commission	
Type	Date	LB	UB	LB	UB
Annuity	Feb 29, 2000	6.07	6.18	5.85	5.98
	Feb 15, 2001	5.56	5.70	5.35	5.47
	June 1, 2001	5.61	5.72	5.40	5.51
	July 2, 2001	5.65	5.76	5.43	5.55
Linear	Feb 29, 2000	6.01	6.18	5.78	5.98
	Feb 15, 2001	5.53	5.70	5.30	5.47
	June 1, 2001	5.57	5.72	5.34	5.51
	July 2, 2001	5.60	5.76	5.36	5.55

full prepayment boundary.

Table 6.2 provides fair rate results for ten year fixed rate periods. Prepayment is restricted to the bottom 22 nodes (per period) of the original lattice and the corresponding nodes in all sublattices, as long as these are located below the full prepayment boundary. This prepayment strategy restricts the number of sublattices. The difference between lower and upper bound can rise up to 20 basis points, although the lower bound is considerably improved compared to the initial lower bound, that is, the fair rate of a non-callable mortgage.

6.8 Concluding remarks

In this chapter a linear programming formulation has been introduced for the valuation and optimal prepayment of (partially) callable mortgages. We have also derived optimal prepayment conditions for fully callable mortgage contracts based on state prices and following from duality theory.

A fully callable mortgage can be modelled by a lattice approach. Partially callable annuity and linear mortgages can only be priced to optimality by an inefficient non-recombining tree approach. To enhance efficiency, we propose a lattice based method to obtain a close lower bound on the fair rate of these mortgage types.

Since, for upward sloping term structures, the fair rate of a partially callable interest-only mortgage (priced to optimality in the chapter 5) provides a practical upper bound on the fair rate of a partially callable annuity, a narrow range for the optimal fair rate is derived. This indicates that the lower bound heuristic is accurate. The comparison of fair rates of non-callable, partially callable and fully callable mortgages is performed extensively in the next chapter.

Related to the LP formulation provided in this chapter, we propose two directions for future research on the optimal valuation of partially callable annuities. First, a theoretical upper bound on the fair rate can be derived by improving the basic dual feasible solution, represented by the full prepayment strategy. The upper bound can be improved by increasing z-variables corresponding to low interest rate states. For these states, an increase in z implies an increase in v_0 by backward recursion, such that the dual objective (mortgage price) increases and the fair rate decreases. The number of z-variables is exponential and therefore many z-variables must be increased from zero to an (a priori unknown) positive value to achieve a significant improvement.

A second direction for further research is based on approximating the fair rate of a partially callable annuity. Since not all states in a non-recombining tree can be included, we might consider a tree defined on a subset of scenario paths. Valuation based on this subtree generates approximative mortgage prices. Approximations are more accurate for finer subtrees. However, approximations can lead to both higher and lower fair rates than the optimal. As a consequence, measuring the accuracy of the approximative fair rate is not possible without the use of fair rate boundaries derived in this chapter.

Chapter 7

A Comparison of Fair Mortgage Rates[1]

7.1 Introduction

Previous chapters have analyzed the valuation of various mortgage contracts. In this chapter we combine results in terms of fair rates. The fair rate is the contract rate at which the mortgage value equals the nominal loan value. If a mortgage is quoted at the fair rate neither bank nor client can make a profit. Fair rate differences between mortgage contracts indicate the value of embedded options. We distinguish annuity, linear and interest-only mortgages. Valuation of savings and investment mortgages is similar to interest-only mortgages, since these mortgage types have a similar cash flow pattern. Various fixed rate periods are considered. Fair rates of mortgages excluding commission costs are stated and we study the effect of including a 1% commission on the fair rate.

Fair rates of non-callable, partially callable and fully callable mortgages are compared. Prepayment is profitable when interest rates are low. However, we have concluded that an optimal prepayment strategy for partially callable mortgages might differ substantially from a full prepayment strategy. Partial prepayment can occur both sooner and later than full prepayment. An early prepayment may occur in December to decrease the unpaid balance before expiration of a prepayment option. A prepayment can be postponed (only for contracts with periodic amortization) to decrease the unpaid balance, since regular redemption payments are much larger before than after an additional prepayment.

[1]This chapter is based on joint work with P. Schotman.

We examine the influence of the underlying term structure model on fair rates, by analyzing mortgage rates based on a one- and two-factor BDT model, applying a square root smoothing function for interest rate volatilities. The underlying interest rate lattices are calibrated on a yield curve based on swap rates and swaption prices, obtained using observed implied Black volatilities. The exact specification and calibration of term structure models has been discussed in chapter 3. Additionally, a HL based model is applied to obtain fair mortgage rates, but a different term structure model did not alter the fair rates significantly. Fair rate results based on the HL model are therefore not reported separately.

All mortgages have a lifetime of 30 years. Interest rate lattices with monthly periods are considered for fair rate calculations. For some instances and mortgage types semi-monthly periods have been applied as well, but fair contract rates never differed by more than 2 basis points compared to the fair rates corresponding to an interest rate lattice with monthly periods. In the remaining of this chapter, fair rates based on monthly periods are reported.

7.2 Results

Fair contract rates for several mortgage types and underlying interest rate models are provided in tables 7.1-7.5, both excluding commission and including a 1% commission. All the examined mortgages have a lifetime of 30 years and a reset date after 5 or 10 years, which can be viewed as a fixed rate period or the first period of an unrestricted ARM. Both non-callable, partially callable and fully callable mortgages are considered. Also, a reconsideration option is included, which is equivalent to a full prepayment option during the last two years of the fixed rate period. All contract rates are quoted as annual rates. Monthly contract rates equal the listed fair rates divided by 12.

Fair rates on variable rate mortgages equal the 1-month short rate and are adapted to the future 1-month rate each period. The difference between the observed variable mortgage rate and the fair rate can be viewed as a premium to cover administration costs plus the significant risk of a variable rate contract. Both client and bank are exposed to the risk of contract rate fluctuations. A change in mortgage rate is immediately reflected, such that each mortgage rate increase is favorable for the bank and each mortgage rate decrease is favorable for the client.

As in the previous chapters we consider fair rate results on four dates. All instances face an upward sloping term structure. A direct result from an increasing term structure is a decreasing contract rate when the fixed rate period shortens. In the extreme case of a variable contract rate the fair rate is lowest. The decrease is largest for the term structure of February 29, 2000, for which the upward slope of the yield curve is steepest.

Tables 7.1-7.3 state fair rate results based on a one-factor BDT interest rate model. Table 7.1 reports fair rates for annuity, linear and interest-only mortgages with a fixed rate period of 10 years, both excluding commission and including a 1% commission. Excluding commission, a full prepayment option is worth 70 to 80 basis points, as the fair rates for fully callable mortgages exceed the fair rates for non-callable mortgages by this amount. A reconsideration option, having the possibility to start a new fixed rate period at the lowest mortgage rate in the last two years of the current fixed rate period, has only a value of 5 basis points. Hence, restricting full prepayment to the final two years decreases the option value dramatically.

A full prepayment option restricted to the initial year(s) is much more valuable. The fair rate of a loan for which full prepayment is restricted to the first two years is on average only 9 basis points less than the fair rate of a fully callable mortgage, whereas a mortgage including a prepayment option for the first year is 16 basis points less.[2] The explanation for this is that the gain of an early prepayment is earned over a longer horizon, whereas a late prepayment is only profitable in a distant future and during a short time span. The value of an early prepayment is not related to the choice of the term structure model or the specification of the volatility function. One may argue that decreasing volatility functions (because of mean reversion) lead to insignificant prepayment decisions due to a non-expanding interest rate lattice for large t, a topic which has been discussed in chapters 2 and 3. This effect would also lead to an early prepayment being much more valuable than a distant future prepayment. However, when using a stationary term structure model (in this case the BDT specification with constant volatility, that is, no mean reversion), the values of full and partial prepayment options do not differ significantly.

As can be concluded from the results, the (regular) amortization schedule which repays the loan as soon as possible gives lowest fair rates. Fair rates are largest for the interest-only mortgages, having the total redemption amount at the end of the contract, middle for the annuities, having increasing redemption amounts over time, and smallest for the linear mortgages, having constant redemption payments. This is consistent with the data,

[2]Note the difference with an entering rate option, which fair rate is equal to the one-period interest rate and is adjusted each period until the fixed-rate period starts.

TABLE 7.1: Fair rates 10 year fixed, one-factor model.

The table reports annualized fair rates for 10 year fixed rate annuities, linear mortgages and interest-only mortgages with a 30 year lifetime. The underlying interest rate process is based on a one-factor BDT model and monthly periods. Interest rate data on four different dates are considered. Included are non-callable mortgages (NC), reconsideration options (REC, mortgage is fully callable in the last two years of the fixed rate period), partially callable mortgages (PC, allowing 20% prepayment each calendar year) and fully callable (FC) mortgages. Contract specifications either exclude commission or include a 1% commission.

Type	Date	NC	REC	PC	FC
No commission					
Annuity	Feb 29, 2000	5.82	5.87	6.07*	6.51
	Feb 15, 2001	5.28	5.34	5.56*	6.09
	June 1, 2001	5.39	5.43	5.61*	6.08
	July 2, 2001	5.41	5.45	5.65*	6.13
Linear	Feb 29, 2000	5.77	5.81	6.01*	6.41
	Feb 15, 2001	5.25	5.31	5.53*	6.02
	June 1, 2001	5.35	5.38	5.57*	6.00
	July 2, 2001	5.37	5.40	5.60*	6.04
Interest-only	Feb 29, 2000	5.86	5.91	6.18	6.57
	Feb 15, 2001	5.30	5.37	5.70	6.14
	June 1, 2001	5.42	5.47	5.72	6.13
	July 2, 2001	5.45	5.50	5.76	6.19
1% commission					
Annuity	Feb 29, 2000	5.68	5.72	5.85*	6.07
	Feb 15, 2001	5.14	5.20	5.35*	5.58
	June 1, 2001	5.25	5.28	5.40*	5.59
	July 2, 2001	5.27	5.31	5.43*	5.64
Linear	Feb 29, 2000	5.62	5.65	5.78*	5.96
	Feb 15, 2001	5.10	5.15	5.30*	5.50
	June 1, 2001	5.20	5.23	5.34*	5.51
	July 2, 2001	5.22	5.25	5.36*	5.54
Interest-only	Feb 29, 2000	5.73	5.77	5.98	6.14
	Feb 15, 2001	5.18	5.24	5.47	5.64
	June 1, 2001	5.30	5.33	5.51	5.66
	July 2, 2001	5.32	5.37	5.55	5.72

* Fair rates of partially callable annuity and linear mortgages are lower bounds, based on the full prepayment boundary and maximally allowed prepayment amounts. Recall the discussion in section 6.6 about bounding the fair rate.

as we are dealing with an increasing term structure and redemption should preferably take place when interest rates are low, implying that the duration of linear mortgage types is shortest, followed by the duration of annuities.

The importance of a 20% partial prepayment option embedded in an interest-only mortgage has been discussed in chapter 5. The values of partial prepayment options in annuity or linear mortgages are similar. This can be concluded from the lower bounds on the fair rate, obtained by the heuristic method described in chapter 6. Note that an upper bound on the fair rate of a traditional mortgage is provided by the fair rate of the interest-only loan, since all term structures are upward sloping. The lower bound is therefore very close to the optimal fair rate, even for a ten year lifetime, for which the heuristic prepayment strategy differs most from the optimal.

Fair rates have also been determined for mortgage loans including a 1% commission. Concerning the ten year fixed rate period, the full prepayment premium is around 40 basis points. The value of a partial prepayment option has significantly increased. An option to prepay 20% per calendar year is worth around 25 basis points or more than 60% of the value of a full prepayment option.

The value of a reconsideration option is still 5 basis points. Also, the value of a full prepayment option restricted to the first years of the fixed rate period has not changed much. If prepayment is restricted to the first two years the option is worth 8 basis points less than a full prepayment option, if prepayment is restricted to the first year the option is still worth 16 basis points less. When comparing results to the case without commission, the compensation for a 1% commission ranges from 15 basis points for non-callable mortgages to 50 basis points for fully callable mortgages.

Table 7.2 shows fair rates for mortgage contracts with a five year fixed rate period, or equivalently, an unrestricted adjustable rate mortgage with an adjustment date after five years. All mortgage rates are lower than the corresponding rates of mortgages with a ten year fixed rate period, since all term structures are upward sloping. In the case without commission costs, the full prepayment premium is about 60 basis points. A reconsideration option valid in the last two years of the fixed rate period (such that the reconsideration period starts after three years) is worth 11 basis points on average. A 20% prepayment option on a five year fixed rate period is worth 19 basis points on average. Partial prepayment options can be priced very accurately, since the lower bound on the fair rate differs not more than 3 basis points from the upper bound, being the fair rate of the corresponding interest-only mortgage.

TABLE 7.2: Fair rates 5 year fixed, one-factor model.

The table reports annualized fair rates for 5 year fixed rate annuities, linear mortgages and interest-only mortgages with a 30 year lifetime. The underlying interest rate process is based on a one-factor BDT model and monthly periods. Interest rate data on four different dates are considered. Included are non-callable mortgages (NC), reconsideration options (REC, mortgage is fully callable in the last two years of the fixed rate period), partially callable mortgages (PC, allowing 20% prepayment each calendar year) and fully callable (FC) mortgages. Contract specifications either exclude commission or include a 1% commission.

Type	Date	NC	REC	PC	FC
No commission					
Annuity	Feb 29, 2000	5.31	5.42	5.48*	5.84
	Feb 15, 2001	4.95	5.08	5.14*	5.59
	June 1, 2001	4.91	5.01	5.10*	5.52
	July 2, 2001	4.86	4.96	5.05*	5.48
Linear	Feb 29, 2000	5.29	5.39	5.46*	5.80
	Feb 15, 2001	4.94	5.07	5.13*	5.57
	June 1, 2001	4.90	4.99	5.08*	5.50
	July 2, 2001	4.85	4.95	5.03*	5.45
Interest-only	Feb 29, 2000	5.32	5.44	5.50	5.87
	Feb 15, 2001	4.96	5.09	5.17	5.61
	June 1, 2001	4.92	5.02	5.10	5.54
	July 2, 2001	4.87	4.98	5.05	5.50
1% commission					
Annuity	Feb 29, 2000	5.07	5.15	5.17*	5.30
	Feb 15, 2001	4.71	4.81	4.82*	4.95
	June 1, 2001	4.68	4.74	4.77*	4.88
	July 2, 2001	4.63	4.70	4.72*	4.84
Linear	Feb 29, 2000	5.04	5.11	5.14*	5.25
	Feb 15, 2001	4.70	4.78	4.79*	4.92
	June 1, 2001	4.66	4.71	4.74*	4.85
	July 2, 2001	4.61	4.67	4.69*	4.80
Interest-only	Feb 29, 2000	5.10	5.18	5.21	5.33
	Feb 15, 2001	4.73	4.83	4.85	4.98
	June 1, 2001	4.70	4.77	4.80	4.91
	July 2, 2001	4.65	4.72	4.75	4.87

* Fair rates of partially callable annuity and linear mortgages are lower bounds, based on the full prepayment boundary and maximally allowed prepayment amounts.

When including 1% commission costs, the average full prepayment premium equals 22 basis points. A 20% partial prepayment option has an average value of 10 basis points. The reconsideration option is worth only 7 basis points on average. For a shorter fixed rate period, the compensation for a 1% commission increases to 25 basis points for non-callable mortgages and 60 basis points for fully callable loans.

Consider now a one-sided bandwidth mortgage, for which an interest rate decrease is immediately reflected in the mortgage rate. In that sense, the bandwidth mortgage is an ARM for which every period is an adjustment date. However, an interest rate increase is reflected only if the increase exceeds a bandwidth, which is 1.25% on the annualized contract rate. If the increase is larger the bandwidth is subtracted from the new mortgage rate. A bandwidth mortgage protects the client from large contract rate increases. Any prepayment option on a bandwidth mortgage is worthless since every interest rate decrease lowers the mortgage rate to its fair rate. The fair rates for non-callable, partially callable and fully callable mortgages are equal. Fair rates for the four instances are 3.94%, 4.85%, 4.71% and 4.66% respectively, which is only slightly higher than the variable mortgage rates (3.49%, 4.71%, 4.64% and 4.64%).

While in the no-commission case a mortgage is called at any interest rate decrease, this is not the case when a commission has to be paid. Then calling takes place only after the commission costs are compensated for by the decrease of the interest rate. The initial contract rate is decreasing faster for mortgage contracts with commission (compared to contracts without commission) when shortening the length of the first fixed rate period. For shorter fixed rate periods commission costs must be compensated for in a shorter time span, by means of a lower fair rate. After the first period we assume that the contract rate can be adjusted without commission or penalty. As a consequence, for very small first fixed rate periods it might not even be possible to find a fair positive contract rate, because the 1% commission costs cannot be compensated by a lower (and still positive) contract rate. For this reason no fair rate for a bandwidth mortgage including commission can be computed, unless the bandwidth is extremely large. In that case the mortgage behaves like a fixed rate mortgage. Calculating fair rates for variable rate mortgages including commission is impossible for the same reason.

Concerning the prepayment option in a mortgage contract, we conclude that this option becomes less valuable when adjustment opportunities increase. Calling a mortgage is less likely when the contract rate is adjusted to the fair rate more often. As an extreme situation, note that the bandwidth mortgage (including only a bandwidth on contract rate

TABLE 7.3: Fair rates two-sided bandwidth, one-factor model, no commission.

The table reports annualized fair rates for annuities, linear mortgages and interest-only mortgages with a 30 year lifetime and an adjustment date after 10 years, including a two-sided bandwidth of 1.25%. The underlying interest rate process is based on a one-factor BDT model and monthly periods. Interest rate data on four different dates are considered. Included are non-callable mortgages (NC), reconsideration options (REC, mortgage is fully callable in the last two years of the fixed rate period) and fully callable (FC) mortgages. All contract specifications exclude commission.

Type	Date	NC	REC	FC
Annuity	Feb 29, 2000	3.91	3.91	3.94
	Feb 15, 2001	4.64	4.64	4.85
	June 1, 2001	4.46	4.46	4.71
	July 2, 2001	4.40	4.40	4.66
Linear	Feb 29, 2000	3.91	3.91	3.94
	Feb 15, 2001	4.64	4.64	4.85
	June 1, 2001	4.46	4.46	4.71
	July 2, 2001	4.40	4.40	4.66
Interest-only	Feb 29, 2000	3.92	3.92	3.94
	Feb 15, 2001	4.64	4.64	4.85
	June 1, 2001	4.46	4.46	4.71
	July 2, 2001	4.40	4.40	4.66

increases) will never be called, as the contract rate is adjusted to the fair contract rate as soon as this fair rate is lower. Hence, there is no reason to prepay the loan before maturity.

Results for two-sided bandwidth mortgages are provided in table 7.3. The bandwidth equals 1.25% on both an interest rate increase and a decrease, protecting both client and bank from (small) interest rate movements. Now the full prepayment premium is positive and can be as large as 25 basis points. The fair rate of a fully callable mortgage with a two-sided bandwidth equals the fair rate of a one-sided bandwidth mortgage, since with a two-sided bandwidth an interest rate decrease does not lead to a contract rate adjustment; instead the mortgage is called, affecting the fair rate similarly.

Comparing tables 7.4 and 7.5 to 7.1 and 7.2, results for the two-factor interest rate model are very close to the fair rates obtained with the one-factor model. Non-callable

TABLE 7.4: Fair rates 10 year fixed, two-factor model.

The table reports annualized fair rates for 10 year fixed rate annuities, linear mortgages and interest-only mortgages with a 30 year lifetime. The underlying interest rate process is based on a two-factor BDT model and monthly periods. Interest rate data on four different dates are considered. Included are non-callable mortgages (NC), reconsideration options (REC, mortgage is fully callable in the last two years of the fixed rate period) and fully callable (FC) mortgages. Contract specifications either exclude commission or include a 1% commission.

Type	Date	NC	REC	FC
No commission				
Annuity	Feb 29, 2000	5.82	5.87	6.49
	Feb 15, 2001	5.28	5.34	6.07
	June 1, 2001	5.39	5.42	6.07
	July 2, 2001	5.41	5.45	6.10
Linear	Feb 29, 2000	5.77	5.81	6.39
	Feb 15, 2001	5.25	5.31	6.01
	June 1, 2001	5.35	5.38	6.00
	July 2, 2001	5.37	5.40	6.02
Interest-only	Feb 29, 2000	5.86	5.91	6.55
	Feb 15, 2001	5.30	5.37	6.13
	June 1, 2001	5.42	5.46	6.13
	July 2, 2001	5.45	5.49	6.16
1% commission				
Annuity	Feb 29, 2000	5.68	5.72	6.05
	Feb 15, 2001	5.14	5.19	5.55
	June 1, 2001	5.25	5.28	5.57
	July 2, 2001	5.27	5.30	5.60
Linear	Feb 29, 2000	5.62	5.65	5.95
	Feb 15, 2001	5.10	5.15	5.47
	June 1, 2001	5.20	5.22	5.49
	July 2, 2001	5.22	5.24	5.51
Interest-only	Feb 29, 2000	5.73	5.77	6.12
	Feb 15, 2001	5.18	5.24	5.62
	June 1, 2001	5.30	5.33	5.64
	July 2, 2001	5.32	5.36	5.67

TABLE 7.5: Fair rates 5 year fixed, two-factor model.

The table reports annualized fair rates for 5 year fixed rate annuities, linear mortgages and interest-only mortgages with a 30 year lifetime. The underlying interest rate process is based on a two-factor BDT model and monthly periods. Interest rate data on four different dates are considered. Included are non-callable mortgages (NC), reconsideration options (REC, mortgage is fully callable in the last two years of the fixed rate period) and fully callable (FC) mortgages. Contract specifications either exclude commission or include a 1% commission.

Type	Date	NC	REC	FC
No commission				
Annuity	Feb 29, 2000	5.31	5.41	5.83
	Feb 15, 2001	4.95	5.08	5.60
	June 1, 2001	4.91	5.01	5.55
	July 2, 2001	4.86	4.96	5.50
Linear	Feb 29, 2000	5.29	5.39	5.78
	Feb 15, 2001	4.94	5.07	5.58
	June 1, 2001	4.90	4.99	5.52
	July 2, 2001	4.85	4.94	5.47
Interest-only	Feb 29, 2000	5.32	5.43	5.85
	Feb 15, 2001	4.96	5.09	5.62
	June 1, 2001	4.92	5.02	5.57
	July 2, 2001	4.87	4.98	5.52
1% commission				
Annuity	Feb 29, 2000	5.07	5.15	5.28
	Feb 15, 2001	4.71	4.80	4.94
	June 1, 2001	4.68	4.74	4.89
	July 2, 2001	4.63	4.69	4.84
Linear	Feb 29, 2000	5.04	5.11	5.23
	Feb 15, 2001	4.70	4.78	4.91
	June 1, 2001	4.66	4.71	4.85
	July 2, 2001	4.61	4.67	4.80
Interest-only	Feb 29, 2000	5.10	5.18	5.32
	Feb 15, 2001	4.73	4.82	4.97
	June 1, 2001	4.70	4.77	4.92
	July 2, 2001	4.65	4.72	4.87

fixed rate contracts are model independent by definition, as only the yield curve (which is matched exactly) affects prices of mortgages without embedded options. Fair rates of fully callable mortgages based on the two-factor model differ from one-factor results 5 basis points incidentally, although the majority of fair rate differences is limited by 2 basis points. Since many fair rates for partially callable mortgages are approximative and a 2 basis point difference is smaller than the approximation range, we have not included results for partially callable mortgages based on an underlying two-factor interest rate model.

A similar robustness holds for the choice of the term structure model. Non-callable fixed rate mortgages only depend on the term structure and are therefore not affected by the underlying term structure model. Fully callable mortgages are only slightly affected. For a ten year fixed rate period, the fair mortgage rate based on the Ho and Lee model is approximately 5 basis points less than the fair rate based on BDT. For shorter fixed rate periods, fair rates are even less sensitive with respect to the underlying term structure model.

Table 7.6 reports both observed[3] mortgage rates and fair rates and provides an indication of a bank's premium to cover costs and risk. The premium on a variable rate mortgage ranges from 0.9 to 1.2 percentage point. Surprisingly, the premium on a five year fixed rate period is larger, at least 1.2 percentage point. A ten year fixed rate period faces an average premium of 1 percentage point. The bandwidth mortgages are very unattractive, since the average premium on a one-sided bandwidth mortgage, with a bandwidth equal to 1.25%, is huge: almost 2 percentage points.

Concerning observed mortgage rates, a contract rate is typically reduced by 30 basis points when 'Nationale Hypotheek Garantie' is included, so even including NHG premiums are large. Also, savings and investment mortgages are usually quoted at a contract rate that is up to 20 basis points higher than annuity, linear or interest-only mortgages. Finally, note that the fair rates for variable rate and bandwidth mortgages are calculated excluding commission, while the observed rates include 1% commission costs in addition.

For the three instances of 2001 the bandwidth mortgage is far less attractive than a fixed rate mortgage. The fair rate is lower than the fair rate of a comparable FRM, while the observed contract rate is higher. For February 29, 2000 the bandwidth mortgage can be attractive as both the real contract rate and the fair rate are significantly lower than the corresponding contract rates of an FRM, although the premium on a bandwidth mortgage is much higher. For clients, the risk of a bandwidth mortgage is high compared

[3] Mortgage rates are observed from Bouwfonds, a Dutch mortgage provider.

TABLE 7.6: Observed rates.

The table provides a comparison between observed and fair contract rates. Mortgage rates charged by banks are observed during the same month in which the term structure of interest rates is rooted. Observed rates exclude 'Nationale Hypotheek Garantie' and are based on a mortgage loan of 100% of the house value. Annuity, linear and interest-only mortgages have similar observed rates. All observed mortgage contracts include 1% commission costs and an allowed prepayment of 20% per calendar year. The corresponding fair rates, stated in brackets, are *upper bounds* based on a one-factor interest rate model, indicating the bank's premium on mortgage loans. Variable rate (VAR), 5 year fixed rate (5Y), 10 year fixed rate (10Y) and 10 year bandwidth (BW) mortgages are considered. Fair rates are excluding commission for variable rate and bandwidth mortgages.

Date	VAR	5Y	10Y	BW
Feb 29, 2000	4.70	6.40	6.90	5.80
	(3.49)	(5.21)	(5.98)	(3.94)
Feb 15, 2001	5.80	6.10	6.50	6.80
	(4.71)	(4.85)	(5.47)	(4.85)
June 1, 2001	5.50	6.00	6.40	6.60
	(4.64)	(4.80)	(5.51)	(4.71)
July 2, 2001	5.50	6.20	6.60	6.70
	(4.64)	(4.75)	(5.55)	(4.66)

to a fixed rate mortgage, especially because the yield curve of February 29, 2000 is very steeply upward sloping. The future contract rate is adjusted continuously, which implies mostly a rate increase (although corrected for the bandwidth). To illustrate the risk of rate adjusting mortgage contracts, figure 7.1 plots the distribution of future contract rates when considering a rate adjustment after 5 years. A contract rate adjustment towards 10% or more has a probability of almost 5%. A contract rate increase is almost inevitable (82%). Note that the initial fair rate equals 5.3%.

Table 7.7 shows the annualized one-month fair rate volatility, given by

$$\sigma_y = 0.5 \cdot (y_u - y_d), \tag{7.1}$$

where y_u and y_d are the annualized fair rates after one up movement or one down movement,

FIGURE 7.1: Distribution of fair future ARM rates.

The figure shows the distribution of fair contract rates of a non-callable adjustable rate annuity (excluding commission) after the adjustment date at 5 years ($t = 60$). The instance considered is of February 29, 2000. The underlying term structure model is a one-factor based BDT model. The vertical dashed line represents the initial contract rate.

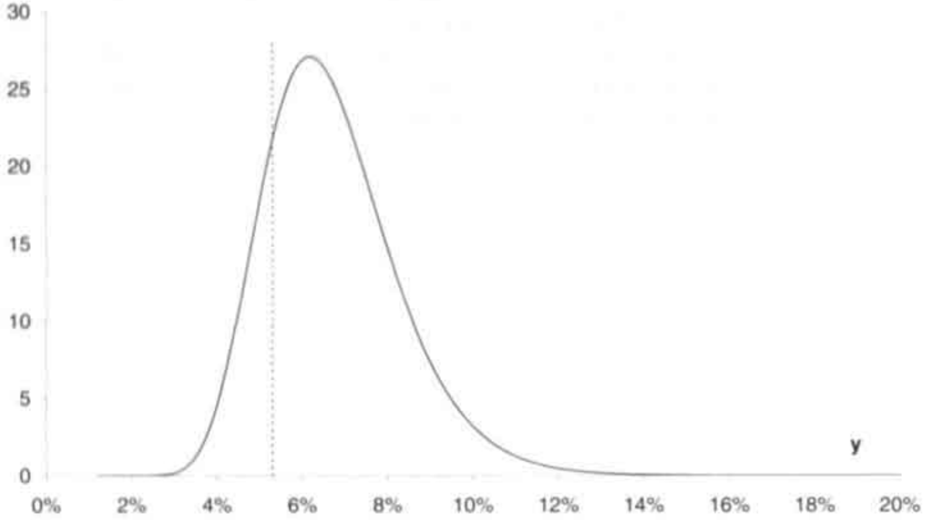

respectively. The maturity date of the fixed rate period remains unchanged. We consider fair rate volatilities of annuity mortgages. Volatilities of linear and interest-only mortgages are similar. The fixed rate period is either five or ten years. Non-callable and fully callable mortgages are included, both excluding commission and including a 1% commission. Fair rate volatilities differ mostly between non-callable and fully callable mortgages. The length of the fixed rate period and the commission do not alter volatilities significantly.

TABLE 7.7: Fair rate volatilities.

The table shows annualized one-month fair rate volatilities in percentages, according to equation 7.1. Non-callable (NC) and fully callable (FC) annuities are considered. All mortgages have a 30 year lifetime. The fixed rate period equals 5 or 10 years. Commission costs are either 0% or 1%. The underlying lattice has monthly periods.

Fixed rate period	10y	10y	10y	10y	5y	5y	5y	5y
Commission (%)	0	0	1	1	0	0	1	1
Call option	NC	FC	NC	FC	NC	FC	NC	FC
Feb 29, 2000	0.207	0.233	0.206	0.229	0.221	0.234	0.220	0.233
Feb 15, 2001	0.175	0.194	0.174	0.193	0.175	0.178	0.174	0.184
June 1, 2001	0.172	0.190	0.172	0.189	0.179	0.182	0.178	0.189
July 2, 2001	0.180	0.195	0.179	0.196	0.180	0.179	0.179	0.189

7.3 Concluding remarks

This chapter provides an overview of fair mortgage rates based on term structures of interest rates on four different dates. The contract rates are robust with respect to the underlying model and the fineness of the grid. Applying a two-factor interest rate model does not change the fair contract rates significantly. Considering 240 semi-monthly periods compared to 120 monthly periods also hardly affects fair rates. Obviously, valuation based on monthly periods is more efficient.

Fair rates for a large variety of mortgage types, including partially callable mortgages analyzed in the previous chapters, are compared. The main results include the value of prepayment options. A full prepayment option is worth 70 to 80 basis points for 10 year fixed rate mortgage contracts excluding commission and 40 basis points for contracts including a 1% commission. The value of a partial prepayment option to prepay 20% per calendar year equals on average 34 basis points excluding commission and 25 basis points including a 1% commission. Consequently, the value of a partial prepayment option embedded in a ten year mortgage contract including commission is more than 60% of a full prepayment option value, although prepayment according to the former option is severely limited.

A reconsideration option, with the possibility to start a new fixed rate period at the lowest contract rate during the last two years of the current period, is only worth 5 to 7 basis points more than a non-callable mortgage. Options restricting prepayment to the first two years of the fixed rate period, fixing the lowest contract rate of the first two years for the remaining lifetime, do have significant value.

Shorter fixed rate periods imply lower fair contract rates, because all term structures considered are upward sloping. Premiums of prepayment options are lower compared to a ten year fixed rate period. An adjustable rate mortgage implies resetting the contract rate to the prevailing future mortgage rate after five years. The lower initial contract rate is offset by an increasing uncertainty about the future contract rate. The future rate is very likely to be higher due to the upward sloping yield curve.

Fair rates have been compared to observed contract rates. For traditional variable and fixed rate mortgage contracts the observed contract rate is roughly 1 percentage point larger than the fair rate. The more complex bandwidth mortgages are expensive, since the premium (between observed and fair rate) on a bandwidth mortgage is almost 2 percentage points.

To complete the discussion on Dutch mortgage contracts, we briefly consider tax regulations in the Netherlands. Interest payments and prepayment penalties are tax deductible, implying that each client will base prepayment decisions on personal (tax and income) circumstances. In a tax free environment the penalty for full prepayment (more than) offsets the gain of full prepayment over the maximally allowed prepayment. Optimally, clients will not prepay more than allowed. Including tax effects, full prepayment is still not optimal. Although a client is partly compensated for the prepayment penalty, future tax deductions decrease because all future interest payments decrease. These partly offsetting arguments introduce a topic for further research on the effect of tax policies on optimal prepayment. As long as prepayment gains and interest payments are equally taxed (by a tax percentage τ), the optimal prepayment strategy will not change and the new fair rate becomes the 'tax free' fair rate multiplied by $\frac{1}{1-\tau}$.

Chapter 8

Summary and Concluding Remarks

This dissertation contributes to two main research fields, mortgage valuation and term structure calibration. Concerning mortgage valuation, computational methods are introduced and analyzed to value restricted prepayment options, present in all Dutch mortgage contracts. For both financial institutions and clients, the importance of mortgage valuation has increased due to the large growth of the Dutch mortgage market in the last decade. For mortgage issuers, one of the largest uncertainties in mortgage contracts concerns prepayment risk. American mortgage loans allow for unrestricted and penalty-free prepayment at any time. Dutch mortgage loans bear less prepayment risk, since only a limited prepayment is allowed penalty-free per calendar year. This so-called partial prepayment option complicates mortgage valuation significantly. Part II of this thesis deals with the valuation of Dutch mortgages.

The second main theme concerns the contribution to the literature on term structure calibration. For pricing interest rate derivatives, including mortgage contracts with embedded options, a term structure of interest rates and a volatility structure are essential. Calibrating a term structure model to a recombining scenario tree (a lattice), in order to match market prices of interest rate derivatives, is discussed in the first part of this thesis.

Performance of term structure models is measured by the ability to price interest rate derivatives accurately. Derivatives we consider include swaps and swaptions. Swap data are available as annual swap rates, which are used to derive a yield curve. Implied swaption volatilities are observed and transformed into prices by Black's formula. Both cash flow patterns and quoting conventions of swaps and swaptions are examined in chapter 2 to obtain a term structure of interest rates and swaption prices.

Chapter 2 also provides an overview of commonly used term structure models, as well as a comparison with respect to swap and swaption pricing based on both model properties and empirical performance. Model prices are compared to prices obtained from observed swap rates and implied swaption volatilities. For our purpose of pricing mortgage contracts and embedded options, a term structure model is selected that is easily calibrated to an interest rate lattice. Empirically desirable properties of term structure models include lognormality and mean reversion of one-period interest rates. Since we consider monthly periods, a model with a limited number of factors is preferred for efficiency reasons.

Calibration of interest rate lattices is based on the Black, Derman and Toy [9, BDT] model. The BDT model, originally defined in discrete time, is easily calibrated to a lattice and captures mean reversion and lognormally distributed interest rates. The original BDT model is a one-factor model. To increase flexibility we also consider a two-factor version.

Besides a detailed analysis of BDT model properties, chapter 3 also provides extensive calibration results. Zero-coupon bond prices are exactly matched by specifying all drift parameters. Swaption pricing errors show a particular pattern. Long term options on short term swaps are underpriced by the model, short term options on long term swaps are overpriced. Average pricing errors are typically smaller than the bid-ask spread for swaptions. Volatilities are hump-shaped. Mean reversion usually starts during the second year, while interest rates are diffusing in the first year. Results are robust, since one- and two-factor models with varying specifications are used for four different dates. Empirically, including multiple factors does not improve calibration results significantly.

Mortgage valuation is the topic of part II of this thesis. The mortgage value equals the present value of all cash flows (redemption, interest payments and additional prepayments). We particularly focus on the valuation of the partial prepayment option. Clients are assumed to exercise prepayment options optimally, based on the development of interest rates. Interest rate scenarios derived in part I serve as input for mortgage valuation.

American mortgage types, allowing full prepayment, can be priced using lattices. Valuation of these mortgage contracts is described in chapter 4. The straightforward pricing method for partially callable mortgages is based on a non-recombining tree approach. Due to the inefficiency of non-recombining trees, we solve mortgage valuation problems including partial prepayment options by applying extended lattice methods.

Some partially callable mortgages can be priced efficiently to optimality. Interest-only mortgages, having no regular periodical amortization, can be viewed as a portfolio of callable bonds. Valuation is based on successively exercising callable bonds, where only one bond can be exercised each calendar year. The portfolio of callable bonds can be valued optimally by an efficient lattice approach, according to chapter 5.

A bond portfolio cannot be used to price partially callable mortgages including a regular amortization schedule (for instance annuities). For this reason, optimal valuation of partially callable annuity mortgages is not possible using a lattice approach. If a non-recombining tree method is applied, optimal prepayment strategies can be derived, but such method is only possible for very small instances. In chapter 6 we formulate a linear programming formulation based on a non-recombining tree. Duality theory and complementary slackness conditions can be applied to derive the optimal prepayment strategy.

A 'no prepayment' strategy is a feasible solution of the primal LP, formulating the valuation of a partially callable mortgage. A 'full prepayment' strategy is a feasible solution of the dual LP. Every feasible solution of the primal LP provides an upper bound on the price of a partially callable mortgage. Similarly, every dual feasible solution provides a lower bound on the price. To obtain a close approximation of the optimal mortgage price (that is, the mortgage price corresponding to an optimal and allowed prepayment strategy), both feasible solutions must be improved in order to narrow the range for the optimal price.

To find a close upper bound on the mortgage price (or equivalently, a lower bound on the fair rate), we propose a heuristic for a prepayment strategy that is close to optimal. An upper bound on the fair rate is obtained by the fair rate of an interest-only mortgage, in case term structures are upward sloping. Combining all results, close to optimal fair rates can be computed efficiently. Directions for further research include the derivation of a theoretical upper bound on the fair rate and the approximation of fair rates by considering a subset of scenario paths.

Chapter 7 concludes the mortgage valuation part. Fair rate results are compared, indicating the values of prepayment options. As an example, for a ten year fixed rate period, a 20% prepayment option is worth more than half the value of a full prepayment option. Additionally, the effect of the yield curve and the fixed rate period on fair rates is discussed, as well as values of contract rate adjustment options and the duration effect of cash flows on the contract rate. Fair rates are robust with respect to the underlying term structure model and the step size of the underlying grid.

An important direction for future research is the effect of the Dutch tax regime on mortgage valuation. Both interest payments and prepayment penalty are tax deductive. Although including tax effects requires a client specific approach, optimal prepayment will not change if a tax-adjusted fair rate is considered.

Bibliography

[1] B. Alink. *Mortgage prepayments in the Netherlands*. PhD thesis, Universiteit Twente, 2002.

[2] K. Amin. On the computation of continuous time option prices using discrete approximations. *Journal of Financial and Quantitative Analysis*, 26(4), 1991.

[3] L Andersen and J. Andreasen. Volatility skews and extensions of the LIBOR market model. *Applied Mathematical Finance*, 7, 2000.

[4] D. Backus, S. Foresi, and S. Zin. Arbitrage opportunities in arbitrage-free models of bond pricing. *Journal of Business and Economic Statistics*, 16(1), 1998.

[5] T. Bali and A. Karagozoglu. Pricing eurodollar futures options using the BDT term structure model: the effect of yield curve smoothing. *Journal of Futures Markets*, 20(3), 2000.

[6] W. Bams. *The term structure of interest rates, a panel data analysis*. PhD thesis, University of Maastricht, 1999.

[7] S. Benninga and Z. Wiener. Binomial term structure models. *Mathematica in education and research*, 7(3), 1998.

[8] F. Black. The pricing of commodity contracts. *Journal of Financial Economics*, 3, 1976.

[9] F. Black, E. Derman, and W. Toy. A one-factor model of interest rates and its application in treasury bond options. *Financial Analysts Journal*, 46(1), 1990.

[10] F. Black and P. Karasinski. Bond and option pricing when short rates are lognormal. *Financial Analysts Journal*, 47(4), 1991.

[11] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81, 1973.

[12] P. Boyle, K. Tan, and W. Tian. Calibrating the Black-Derman-Toy model: some theoretical results. *Applied Mathematical Finance*, 8(1), 2001.

[13] A. Brace, D. Gatarek, and M. Musiela. The market model of interest rate dynamics. *Mathematical Finance*, 7(2), 1997.

[14] M. Brennan and E. Schwartz. An equilibrium model of bond pricing and a test of market efficiency. *Journal of Financial and Quantitative Analysis*, 17(3), 1982.

[15] D. Brigo and F. Mercurio. *Interest rate models: theory and practice*. Springer Finance, 2001.

[16] W. Bühler, M. Uhrig-Homburg, U. Walter, and T. Weber. An empirical comparison of forward-rate and spot-rate models for valuing interest rate options. *Journal of Finance*, 54(1), 1999.

[17] A. van Bussel. *Valuation and interest rate risk of mortgages in the Netherlands*. PhD thesis, University of Maastricht, 1998.

[18] K. Chan, G. Karolyi, F. Longstaff, and A. Sanders. An empirical comparison of alternative models of the short-term interest rate. *Journal of Finance*, 47(3), 1992.

[19] E. Charlier and A. Van Bussel. Prepayment behavior of Dutch mortgagors: An empirical analysis. *Real Estate Economics*, 31(2), 2003.

[20] L. Chen. *Interest rate dynamics, derivatives pricing and risk management*. Springer, 1996.

[21] T. Cheuk and T. Vorst. Shout floors. *NetExposure - Electronic Journal of Financial Risk*, 2, 1997.

[22] V. Chvátal. *Linear programming*. W.H. Freeman and Company, 1983.

[23] J. Cox, J. Ingersoll, and S. Ross. A theory of the term structure of interest rates. *Econometrica*, 53, 1985.

[24] F. De Jong, J. Driessen, and A. Pelsser. On the information in the interest rate term structure and option prices. *Review of Derivatives Research*, 7(2), 2004.

[25] E. Derman and I. Kani. Riding on a smile. *RISK*, 7(2), 1994.

[26] J. Driessen, P. Klaassen, and B. Melenberg. The performance of multi-factor term structure models for pricing and hedging caps and swaptions. *Journal of Financial and Quantitative Analysis*, 38(3), 2003.

[27] D. Duffie. *Dynamic asset pricing theory*. Princeton University Press, 2^{nd} edition, 1996.

[28] D. Duffie and K. Singleton. An econometric model of the term structure of interest-rate swap yields. *Journal of Finance*, 52(4), 1997.

[29] B. Dumas, J. Fleming, and R. Whaley. Implied volatility functions: Empirical tests. *Journal of Finance*, 53(6), 1998.

[30] F. Fabozzi. *Investment management*. Prentice Hall College, 2^{nd} edition, 1999.

[31] R. Fan, A. Gupta, and P. Ritchken. On the performance of multi-factor term structure models for pricing caps and swaptions. *Working paper, Case Western Reserve University*, 2001.

[32] L. Hayre. Prepayment modeling and valuation of Dutch mortgages. *Journal of Fixed Income*, 12(4), 2003.

[33] D. Heath, R. Jarrow, and A. Morton. Bond pricing and the term structure of interest rates: a discrete time approximation. *Journal of Financial and Quantitative Analysis*, 25(4), 1990.

[34] D. Heath, R. Jarrow, and A. Morton. Bond pricing and the term structure of interest rates: a new methodology. *Econometrica*, 60(1), 1992.

[35] P. Hendershott and R. Van Order. Pricing mortgages: An interpretation of the models and results. *Journal of Financial Services Research*, 1(1), 1987.

[36] J. Hilliard, J. Kau, D. Keenan, and W. Muller III. Pricing a class of American and European path dependent securities. *Management science*, 41(12), 1995.

[37] J. Hilliard, J. Kau, and V. Slawson Jr. Valuing prepayment and default in a fixed-rate mortgage: a bivariate binomial options pricing technique. *Real Estate Economics*, 26(3), 1998.

[38] T. Ho and S. Lee. Term structure movements and pricing interest rate contingent claims. *Journal of Finance*, 41(5), 1986.

[39] J. Hull. *Options, futures, and other derivatives*. Pearson Education, 5^{th} edition, 2002.

[40] J. Hull and A. White. The pricing of options on assets with stochastic volatility. *Journal of Finance*, 42(2), 1987.

[41] J. Hull and A. White. Pricing interest rate derivative securities. *Review of Financial Studies*, 3(4), 1990.

[42] J. Hull and A. White. One-factor interest rate models and the valuation of interest rate derivative securities. *Journal of Financial and Quantitative Analysis*, 28(2), 1993.

[43] J. Ingersoll Jr. *Theory of financial decision making*. Rowman and Littlefield, 1987.

[44] J. James and N. Webber. *Interest rate modelling*. John Wiley & Sons, 2000.

[45] F. Jamshidian. LIBOR and swap market models and measures. *Finance and Stochastics*, 1(4), 1997.

[46] R. Jarrow. *Modelling fixed income securities and interest rate options*. McGraw-Hill, 1996.

[47] J. Kau and D. Keenan. An overview of the option-theoretic pricing of mortgages. *Journal of Housing Research*, 6(2), 1995.

[48] J. Kau, D. Keenan, W. Muller III, and J. Epperson. The valuation and analysis of adjustable rate mortgages. *Management Science*, 36(12), 1990.

[49] J. Kau, D. Keenan, W. Muller III, and J. Epperson. A general valuation model for fixed-rate residential mortgages. *Journal of Money, Credit and Banking*, 24(3), 1992.

[50] J. Kau, D. Keenan, W. Muller III, and J. Epperson. Option theory and floating-rate securities with a comparison of adjustable- and fixed-rate mortgages. *Journal of Business*, 66(4), 1993.

[51] J. Kau, D. Keenan, W. Muller III, and J. Epperson. The valuation at origination of fixed rate mortgages with default and prepayment. *Journal of Real Estate Finance and Economics*, 11(1), 1998.

[52] J. Kau and V. Slawson Jr. Frictions, heterogeneity and optimality in mortgage modelling. *Journal of Real Estate Finance and Economics*, 24(3), 2002.

[53] A. Kolen. Determining the fair value of a mortgage with partial prepayment: A mathematical programming approach. *Research memorandum University of Maastricht*, 2003.

[54] W. Leung and C. Sirmans. A lattice approach to pricing fixed-rate mortgages with default and prepayment options. *Journal of the American Real Estate and Urban Economics Association*, 21(3), 1990.

[55] A. Li, P. Ritchken, and L. Sankarasubramanian. Lattice models for pricing American interest rate claims. *Journal of Finance*, 50(2), 1995.

[56] F. Longstaff, P. Santa-Clara, and E. Schwartz. The relative valuation of caps and swaptions: Theory and empirical evidence. *Journal of Finance*, 56(6), 2001.

[57] F. Longstaff and E. Schwartz. Interest rate volatility and the term structure: A two factor general equilibrium model. *Journal of Finance*, 47(4), 1992.

[58] F. Longstaff and E. Schwartz. Valuing American options by simulation: a simple least squares approach. *Review of Financial Studies*, 14(1), 2001.

[59] D. Luenberger. *Investment science*. Oxford University Press, 1998.

[60] Y. Lyuu. *Financial engineering and computation: principles, mathematics, algorithms*. Cambridge University Press, 2002.

[61] R. Mathis III and G. Bierwag. Pricing eurodollar futures options with the Ho and Lee and Black, Derman, Toy models: an empirical comparison. *Journal of Futures Markets*, 19(3), 1999.

[62] J. McConnell and M. Singh. Rational prepayments and the valuation of collateralized mortgage obligations. *Journal of Finance*, 49(3), 1994.

[63] J. McCulloch. The tax adjusted yield curve. *Journal of Finance*, 30(3), 1975.

[64] F. Mercurio and J. Moraleda. A family of humped volatility models. *European Journal of Finance*, 7(2), 2001.

[65] R. Merton. Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4, 1973.

[66] K. Miltersen, K. Sandmann, and D. Sondermann. Closed form solutions for term structure derivatives with log-normal interest rates. *Journal of Finance*, 52(1), 1997.

[67] C. Nelson and A. Siegel. Parsimonious modelling of yield curves. *Journal of Business*, 60(4), 1987.

[68] D. Nelson and K. Ramaswamy. Simple binomial processes as diffusion approximations in financial models. *Review of Financial Studies*, 3(3), 1990.

[69] NEOS-Guide. www-fp.mcs.anl.gov/otc/guide/index.html, 1996.

[70] S. Nielsen and R. Poulsen. A two-factor, stochastic programming model of Danish mortgage-backed securities. *Journal of Economic Dynamics and Control*, 28(7), 2004.

[71] A. Pelsser. *Efficient valuation of interest rate derivatives*. Springer, 2000.

[72] S. Peterson, R. Stapleton, and M. Subrahmanyam. A multi-factor spot-rate model for the pricing of interest-rate derivatives. *Journal of Financial and Quantitative Analysis*, 38(4), 2003.

[73] S. Pliska. *Introduction to mathematical finance: discrete time models*. Blackwell, 1997.

[74] J. Prendergast. The complexities of mortgage options. *Journal of Fixed Income*, 12(4), 2003.

[75] R. Rebonato. *Interest rate option models*. John Wiley & Sons, 2^{nd} edition, 1998.

[76] R. Rebonato. On the simultaneous calibration of multi-factor log-normal interest rate models to Black volatilities and to the correlation matrix. *Journal of Computational Finance*, 2(4), 1999.

[77] P. Ritchken and I. Chuang. Interest rate option pricing with volatility humps. *Review of Derivative Research*, 3(3), 2000.

[78] W. Schmidt. On a general class of one-factor models for the term structure of interest rates. *Finance and Stochastics*, 1(1), 1996.

[79] O. Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5, 1977.

Samenvatting / Summary in Dutch

Hypotheekwaardering en rentemodellen vormen de twee hoofdthema's in deze dissertatie. De belangrijkste bijdrage aan de bestaande literatuur over het prijzen van hypotheken is de waardebepaling van een optie tot gedeeltelijk aflossen, die in alle Nederlandse hypotheken aanwezig is. We veronderstellen dat (een gedeelte van) een hypotheeklening wordt afgelost op het moment dat rentes laag zijn. Om die reden worden rentemodellen geanalyseerd in het eerste deel van dit proefschrift, voorafgaand aan hypotheekwaardering.

Een hypotheek is een lening, verstrekt door een bank of andere financiële instelling, met onroerend goed als onderpand. De bank vervult de rol van hypotheeknemer, de klant stelt het onderpand beschikbaar en is daarmee hypotheekgever. Het onderpand dient als garantie voor de bank als de klant de overeengekomen periodieke betalingen niet nakomt.

Het belang van het waarderen van hypotheken voor alle financiële instellingen neemt de laatste jaren sterk toe. Zowel banken als institutionele beleggers geven nieuwe hypotheekleningen uit of beleggen in bestaande leningen. Het totale bedrag aan uitstaande hypotheken in Nederland is verdrievoudigd in de afgelopen tien jaar. In Europa bezet de Nederlandse hypothekenmarkt de tweede plaats op basis van het bedrag aan uitstaande leningen, ondanks een relatief klein aantal inwoners. De belangrijkste oorzaak hiervoor is het gunstige Nederlandse belastingklimaat, dat hypotheekgevers netto een goedkope mogelijkheid biedt om hoge hypotheeklasten aan te houden.

Vervroegd aflossen van hypotheken vormt een groot risico voor financiële instellingen. Amerikaanse hypotheken laten ongelimiteerde en boetevrije aflossingen toe op elk moment. Nederlandse hypotheken zijn aan minder aflossingsrisico onderhevig, omdat slechts een beperkt gedeelte van de lening boetevrij mag worden afgelost per kalenderjaar. Deze gedeeltelijke aflossingsoptie bemoeilijkt hypotheekwaardering in grote mate.

Voor het waarderen van hypotheekcontracten en bijbehorende (aflossings)opties zijn een rentetermijnstructuur en de volatiliteiten van de rentes essentieel. Beide worden gemodelleerd met behulp van scenario's, die gecombineerd worden tot een renteboom. Een renteboom is een discrete weergave van de continue verdeling van mogelijke rentes in de toekomst. Voor het calibreren van een renteboom maken we gebruik van geobserveerde swaprentes en swaptionprijzen. Een renteboom is gecalibreerd als swaps en swaptions correct geprijsd worden, dat wil zeggen als de modelprijzen overeenkomen met de geobserveerde data.

Hoofdstuk 2 geeft een overzicht van populaire rentemodellen en hun kenmerken. Ook wordt de waardering van swaps en swaptions uitgebreid beschreven, gebaseerd op de kasstromen van deze rentederivaten. Ten slotte geeft dit hoofdstuk een literatuuroverzicht waarin de prestaties van verschillende modellen met betrekking tot het prijzen van swaps en swaptions worden uiteengezet.

Calibratie geschiedt op basis van verschillende variaties op het Black, Derman en Toy [9] rentemodel. Hoofdstuk 3 analyseert de kenmerken van dit model, waaronder 'mean reversion' en lognormaal verdeelde rentes. De gemodelleerde rentetermijnstructuur komt exact overeen met de geobserveerde termijnstructuur gebaseerd op swaprentes. De gemiddelde afwijking tussen geobserveerde prijzen en modelprijzen van swaptions is kleiner dan de bid-ask spread. Swaptions met een lange optielooptijd en een korte swaplooptijd worden door het model ondergeprijsd, swaptions met een korte optielooptijd en een lange swaplooptijd worden overgeprijsd. De volatiliteit van de korte rente heeft een karakteristieke 'hump'. Rentes zijn beperkt 'mean reverting', in het eerste jaar is er zelfs sprake van divergentie.

De gecalibreerde rentebomen worden gebruikt voor het waarderen van hypotheken en voor het bepalen van de optimale aflossingsstrategie. We vergelijken hypotheken op basis van 'eerlijke' contractrentes (fair rates). De hypotheekrente is fair als de som van alle verdisconteerde rentebetalingen en aflossingen exact gelijk is aan de nominale waarde van de lening. Bij deze rente maken zowel klant als bank geen winst, gegeven de verwachte ontwikkeling van rente en volatiliteit. Fair rates worden vooral gebruikt om aflossingsopties te waarderen. Zo geeft het verschil tussen de fair rate van een ongerestricteerd aflosbare hypotheek en die van een niet aflosbare de waarde van een ongerestricteerde aflossingsoptie.

In hoofdstuk 4 onderscheiden we hypotheekcontracten op basis van aflossingsschema, opties tot vervroegd aflossen en opties tot het aanpassen van de contractrente. Hier worden de meest gangbare hypotheekvormen beschreven die (met behulp van een recombinerende scenarioboom) efficient gewaardeerd kunnen worden. De optimale aflossingsstrategie van een onbeperkt aflosbare hypotheeklening komt aan de orde.

Hypotheken met gedeeltelijke aflossingen vereisen complexere waarderingstechnieken.

In hoofdstuk 5 ontwikkelen we een efficient algoritme voor het prijzen van beperkt aflosbare 'interest-only' hypotheken. Deze methode is gebaseerd op het opsplitsen van een hypotheek in een aantal volledig aflosbare obligaties. De in Nederland populaire 'interest-only' hypotheken kunnen met dit algoritme gewaardeerd worden omdat deze contracten geen reguliere maandelijkse aflossing kennen.

Hypotheken met reguliere aflossingen kunnen niet zonder meer opgesplitst worden in obligaties. Waardering is ingewikkelder omdat vervroegde aflossingen ook het reguliere aflossingspatroon beïnvloeden. In hoofdstuk 6 formuleren we een lineair programmeringsmodel (LP), waarin de som van alle verdisconteerde betalingen wordt geminimaliseerd. Op basis van LP dualiteit kan een optimale aflossingsstrategie voor onbeperkt aflosbare hypotheken worden afgeleid. Prijzen van niet en onbeperkt aflosbare hypotheken begrenzen de waarde van een beperkt aflosbare hypotheek. 'Niet aflossen' vormt een toegelaten oplossing van het primale LP en geeft een bovengrens voor de prijs van een beperkt aflosbare hypotheek. 'Onbeperkt aflossen' is een toegelaten oplossing van het duale LP en geeft een ondergrens voor de prijs.

Door het interval tussen bovengrens en ondergrens te verkleinen, wordt een nauwkeurige schatting van de prijs (of van de fair rate) verkregen. Omdat aan het LP een niet-recombinerende boom ten grondslag ligt, met een exponentiële toename van het aantal scenario's, kunnen grote instanties niet tegelijk efficient en optimaal worden opgelost. Om die reden leiden we een suboptimale aflossingsstrategie af, die een goede bovengrens op de prijs oplevert.

Hoofdstuk 7 sluit het tweede deel af met een uitgebreid overzicht van fair rates van de meest voorkomende hypotheken, gecategoriseerd naar aflossingspatroon, renteaanpassingen en toegestane aflossingsmogelijkheden. Hieruit volgen onder andere de waardes van aflossingsopties. Empirisch blijkt dat, voor een aantal typische hypotheken, een optie om 20% van de oorspronkelijke lening per kalenderjaar af te lossen ten minste de helft waard is van een onbeperkt toegestane aflossing.

Ook komen het effect van de rentevastperiode en de termijnstructuur op de fair rate en de invloed van de duratie van kasstromen op de contractrente aan de orde. Waargenomen contractrentes worden vergeleken met fair rates om een indicatie te krijgen van de premie voor banken. Deze ligt voor basishypotheken rond 1 procentpunt, hypotheken met gecompliceerde opties kennen een hogere premie.

Curriculum Vitae

Bart Kuijpers was born on April 24, 1978 in Echt, The Netherlands. In 1995 he graduated cum laude from the Bisschoppelijk College Echt. From 1995 to 1999 he has studied econometrics at the University of Maastricht, from which he received cum laude the degree Master of Science in Operations Research in September 1999. He wrote his master's thesis during an internship at KPN Research in Leidschendam.

From October 1999 till March 2004, Bart Kuijpers was a Ph.D. candidate at the Departments of Finance and Quantitative Economics, at the Faculty of Economics and Business Administration of the University of Maastricht. His research on mortgage valuation and term structure models was supervised by Prof. dr. ir. Antoon Kolen and Prof. dr. Peter Schotman.

In January 2002, Bart Kuijpers received the certificate of the Dutch Network of Mathematical Operations Research (LNMB). He was involved in executive teaching for Asset Liability Management at the University of Maastricht and was active in graduate teaching for the Departments of Finance and Quantitative Economics.

From April 1, 2004, Bart Kuijpers works for Optas Management BV in Rotterdam.