Danish Mortgage-Backed Bonds
Term Structure and Prepayment Modelling

Master's Thesis

Erik Bennike & Peter Rasmussen

Final Version

April 7, 2006

University of Copenhagen
Department of Economics
Advisor: Peter Norman Sørensen

Summary

In Denmark, a typical way of financing acquisition of real estate property is mortgage financing. A mortgage-backed bond (mortgage bond) is a bond secured by a mortgage on real estate property is a bond. What makes Danish mortgage bonds particularly interesting, is that a large part of the outstanding bonds are long-term bonds with an embedded prepayment option. The prepayment option gives the mortgagor (borrower) the right to redeem the loan at par at (almost) any time. Due to the existence of the prepayment option, these bonds are referred to as *callable* mortgage bonds. Had it not been for this embedded prepayment option, a mortgage bond would be fairly easy to price. It would simply be a matter of discounting the scheduled cash flows of the bond with a relevant yield curve, e.g. a government or swap yield curve added a spread for credit risk, liquidity etc. Consequently, non-callable mortgage bonds are not overly difficult to price.

However, callable mortgage bonds are very difficult to price. The prepayment option makes the cash flow of the bond uncertain, since it cannot be known initially when or if the prepayment option will be exercised. It is exactly the uncertainty of the cash flow of a callable mortgage bond that makes it intriguingly interesting to model.

The two main ingredients of a model for pricing callable mortgage bonds are a term structure model and a prepayment model. A term structure model is a stochastic model of the evolution of the term structure. Hence, the term structure model dictates a stochastic pattern for the evolution of the term structure. The reason why we need to model the evolution of the term structure is to make estimates of the size of exercise of the prepayment option in the future. This, in turn, is done to determine a probability weighted cash flow of the bond. When the term structure model is set up, calibrated and implemented, the next step is to to create a prepayment model that estimates the size of prepayments given inputs from among other things, the term structure model.

We start out the thesis by setting up a general pricing framework based on basic assumptions of no arbitrage, frictionless markets etc. We develop a pricing framework based on the martingale approach, where we price assets by replacing the real-world probability measure with the martingale probability measure. In the general pricing framework we aim at developing formulas for the price of a zero-coupon bond, since this will enable us to price also more complicated assets as portfolios of zero coupon bonds. Furthermore, we derive formulas for European options on zero coupon bonds, since these prices will be used to calibrate the term structure model in a later section.

We then turn towards the modelling of the term structure of interest rates, starting

with a practical section on how to actually derive an initial yield curve using observed bond prices. Having done this, we take the modelling of the evolution of the term structure under treatment, starting with a short review of existing models. We find that the Hull-White term structure model is a good choice of model for the purpose at hand. We build on the derived price formulas from the general pricing section to create formulas for European options on zero coupon bonds under the Hull-White model.

However, for the calibration issue, we want to use interest rate caps and floors as calibrating instruments, and we therefore derive formulas for the value of such instruments. This is done by the use of option prices, taking advantage of the observation that a cap can be seen as a collection of put options on zero coupon bonds. Equivalently, a floor can be seen as a collection of call options on zero coupon bonds. We proceed to calibrate the model. We do this by matching observed and model prices of interest rate caps. When we have obtained the estimates of the parameters of the model, we can implement the model.

The Hull-White model is usually implemented using a trinomial interest rate tree, and this is exactly what we do, creating an interest rate tree based on our derived yield curve and estimated parameters. Throughout the section, we put strong emphasis on the practical aspects involved in the exercise along with the theoretical arguments. We create an interest rate tree for the first eight quarters, deriving the martingale probabilities and the interest rates in the interest rate tree.

When we have implemented the term structure model by creating the interest rate tree, we start the treatment of prepayment modelling. We do this by first establishing some arguments of what should drive prepayments in a framework based on an assumption of rational behavior. In general, rational prepayment models dictate the the mortgagor should prepay his mortgage loan, every time the value of the existing debt exceeds the value of the refinancing alternative. Hence, such a model, implies that the price of the mortgage cannot exceed par, if markets are frictionless. We investigate an extension within the rational prepayment set-up, where costs of prepayments are taken into account. Even though this extension improves the rational prepayment set-up by providing the possibility of prices above par and running prepayments, providing a heterogeneity in the loan sizes.

We proceed to describe a few important actual drivers of prepayments. These entail of course the economic gain of prepayment, but also the maturity of the loan and the loan size. The economic gain of prepayment must be expected to be the single most important driver of prepayments. This provides the basis for the modelling of prepayments. We begin by reviewing the proprietary US model of Goldman Sachs and a more recently constructed Danish model, proceeding to create our own prepayment model.

Initially, we set up a model with the economic gain of prepayment, the average loan size, and the time to maturity as explanatory variables. We use a probit formulation, and we estimate the parameters of the model by maximum likelihood. Of the three variables included, we find the average loan size to be statistically insignificant. We therefore exclude this variable from the model, but on the other hand, we include two new variables. These are the slope of the yield curve and the change in the refinancing interest rate. These are both found to be significant drivers of prepayment in the specified model. This model achieves an explanatory power of approximately 72%, which can be said to be satisfactory. Especially when we consider the simplifying assumptions made along the way, the results are very encouraging. We round off the section by discussing possible extensions to the formulated prepayment model. We give special attention to the issue of the applicability of the so-called preliminary (or scheduled) prepayments as a signal of final prepayments at the next term. We find a very solid pattern, and our results indicate that much can be gained in the modelling of prepayments by including preliminary prepayments in the estimation of prepayments at the next term.

We finish the pricing sections of the thesis by providing an overview of how the term structure model and the prepayment model can be combined to calculate the fair value of a callable mortgage bond.

Then we turn our view towards the investment issue. We start the treatment of this by describing relevant return and risk measures, both for callable and non-callable bonds. We present an application, in which we calculate various key figures of callable bonds and a non-callable bond to facilitate the understanding of the differences between these two types of bonds.

In the following section, we go more deeply into the technical aspects of how to set up an investment strategy including callable mortgage bonds. Initially, we explain how to carry out static hedges of interest rate risk for bonds in general, using first and second derivatives of the price-yield relationship. Subsequently, we apply the techniques of hedging in creation of our own trading strategies for mortgage bonds. We create a portfolio of swaptions and a non-callable government bond in order to track the first and second derivatives of the price-yield relationship of a particular callable mortgage bond. This enables us to evaluate the richness of the callable mortgage bond by comparing the holding period return for the tracking portfolio and the callable mortgage bond, for a broad spectrum of parallel shifts in the yield curve. We end the trading strategies section by sketching an example of how to implement a prepayment bet, using a simple example of differences in debtor distributions.

We close the thesis with a short discussion of the product innovation that has been taking place on the Danish mortgage market in recent years. We discuss the challenges it poses to market participants, and sketch the principle of pricing these bonds. In particular, we address the construction of adjustable rate mortgages, capped floating loans and instalment-free loans and their construction on the bond market.

Preface

After both having spent time studying abroad, we went searching for a perfect topic for our Master's thesis during the summer 2005. We drew on good experiences from our Bachelor's project, which we also wrote together, and decided to make a joint thesis. The choice of a topic took some time of considerations, but we decided that mortgage-backed bonds would be a excellent topic, since it gave us the possibility to include three factors that we both found interesting to include in the thesis. These were (1) that the topic should clearly be related to finance, (2) the topic should be treating complicated financial products, and (3) the topic should give the possibility to do some work of our own, both theoretical and empirical modelling. The topic of mortgage-backed bond had the ability to combine these three factors.

We have had easy access to data and analytics packages at our student jobs at HSH Nordbank Copenhagen Branch and Danske Bank, respectively. We have had unlimited access to the data sources and function libraries available here, but have otherwise received little guidance or help with the thesis. We thank both of our employers for providing us with this access and our colleagues for their cheering support.

Furthermore, we are grateful to Martin D. Linderstrøm and Frederik Silbye for providing helpful comments with the final draft. Beyond that, we have not received any help with the thesis, and we therefore of course remain responsible for all errors left in it.

To comply with existing rules, we have chosen to divide the thesis, such that Erik is responsible for sections 3.1, 3.3, 3.4, 4.1, 4.3.2, 4.3.3, 5, 6 and 9, while Peter is responsible for sections 2, 3.2, 4.2, 4.3.1, 7 and 8. We remain jointly responsible for the introduction and the conclusion (sections 1 and 10). However, we strongly emphasize that the entire thesis is to be seen as the result of a *joint* effort, and we encourage the reader to regard it as such.

Copenhagen, April 2006

Erik Bennike & Peter Rasmussen

Contents

1 Introduction

1.1 Mortgage-Backed Bonds

A mortgage-backed bond – or in short a mortgage bond – is a bond secured by a mortgage on real estate property. Mortgage financing is a very common way of financing acquisition of real estate in especially US and Northern Europe, including Denmark.[1] Why should one give special attention to the pricing of these bonds? The answer lies in the intriguing complexity of the product and its widespread application in Denmark.

In Figure 1.1, we show price-yield pairs of a Danish government bond and a mortgage bond, respectively. The prices of these two bonds are plotted with the yield to maturity of the government bond on the abscissa axis.

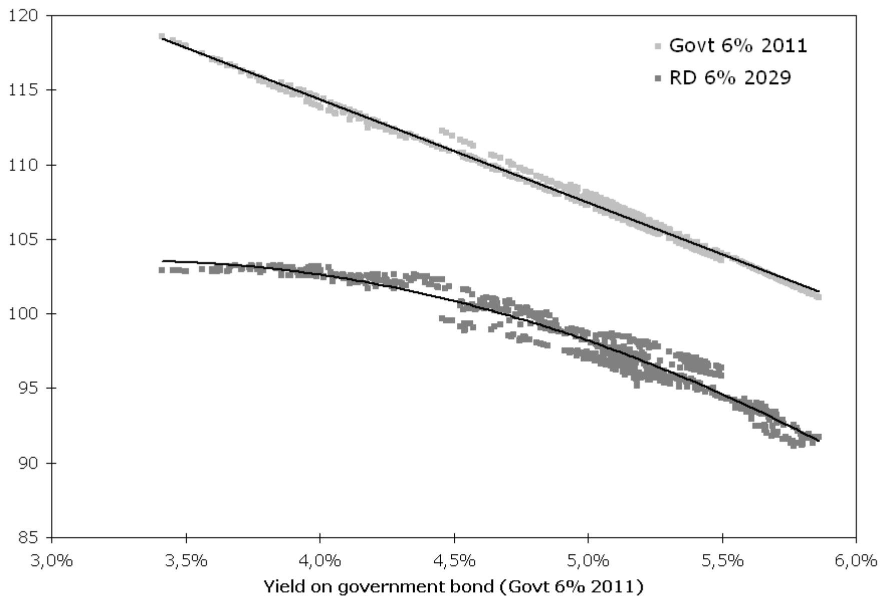

Source: Own calculations based on price data from Copenhagen Stock Exchange.
Note: The data period is June 1, 2000 – January 11, 2006.

Figure 1.1: Prices of Govt 6% 2011 and RD 6% 2029 plotted against yields on Govt 6% 2011

This figure provides clear motivation as to why one should give mortgage bonds special interest. First looking at the government bond, the price-yield relationship of this bond is very standard. This Danish government bond is a bullet bond and has no embedded options or other derivatives of any kind. Therefore, the price-yield relationship is moderately negative and close to linear. However, the mortgage bond is much more interesting. Two things are worth noting. The first one is the negative curvature of the price-yield

[1] In this thesis we focus specifically on the Danish market for mortgage bond products, but we do make references to the US market in particular along the way.

relationship. This is referred to as *negative convexity*. The term covers the fact that the first derivative of the price-yield relationship is a decreasing function of the yield, i.e. the second derivative of the price-yield relationship is negative. We treat this and related issues more in-depth in section 7. The second issue worth noting in Figure 1.1 is the existence of some kind of a price ceiling over the price of the bond. These two issues are special features of Danish mortgage bonds.

The existence of the negative convexity and the price ceiling is both due to the *prepayment option* embedded in these bonds. A prepayment option on a Danish mortgage bond is an option that gives the mortgagor[2] the right to redeem the loan at prespecified quarterly exercise dates along the life of the bond. Mortgage bonds with such an embedded prepayment option are referred to as *callable mortgage bonds*.

The mortgagor has an incentive to exercise the option, i.e. buying the bonds back at par, if the price exceeds par. Most exercises of the prepayment option happens in connection with loan conversion, which means that the mortgagor prepays his loan with a given coupon rate, and takes on different loan with a lower coupon rate. The incentive to do so is obviously high if the coupon rate is significantly higher than the refinancing rate, i.e. the coupon rate that a new mortgage loan will have. Thus, many mortgagors can be expected to prepay their loans if the interest rate falls. Therefore, the price increases of the callable mortgage bond becomes smaller and smaller as the interest rate decreases, since the investors will not be willing to pay so much for the mortgage bond, expecting a high level of prepayments. In the end, for sufficiently low yield levels, the price-yield relationship can actually become positive, as the prepayment incentive becomes extremely high. Hence, the prepayment option gives rise to the negative convexity pattern shown in Figure 1.1.

The pricing set-up of a callable mortgage bond is very different from the pricing set-up of a non-callable government bond or mortgage bond. The pricing of non-callable bonds is relatively straightforward, since the only thing that is needed is actually a relevant yield curve, perhaps with a relevant credit spread. Discounting the cash flow of the non-callable bond according to the yield curve will provide a fair value.

When pricing callable bonds, things start to get extremely complicated. This is all due to the prepayment option. The prepayment option implies the risk of the investment being called soon after the acquisition causing a possible considerable loss, if the bonds have been bought at a price beyond par. We already at this premature stage write the

[2]The borrower in a mortgage loan arrangement.

value of the callable bond as

$$P_{callable} = P_{noncallable} - P_{call\ option} \tag{1.1}$$

Hence, the value of a callable bond is the value of a non-callable bond with similar properties less the value of a call option on the bond. However, pricing the prepayment option is very difficult.

In order to find a fair value of a callable mortgage bond, one needs to obtain an estimate of how many loans that will be prepaid in the future. The level of prepayments at a given point in time is influenced primarily by one factor – the yield curve. Therefore, to estimate the prepayment extent in future periods, a model for the yield curve in the future is needed – a term structure model. Once this is obtained, a model for the prepayments can be applied to the results of the model governing the stochastic interest rate development in the future.

Pricing mortgage bonds is a very complicated issue. The model set-up needed is indeed comprehensive. Duarte, Longstaff & Yu (forthcoming) note that these models require a high level of intellectual capital to develop, maintain and use. Typical investors of these bonds are therefore also commercial banks, pension and insurance funds, mutual funds etc. Such professional investors typically have the capability of setting up pricing models for securities as complex as these, and yet fortunes are spent on this topic in these institutions. The investor distributions of Danish government bonds and mortgage bonds are shown in Figure 1.2.

From this figure, it is also seen that the domestic financial sector, which consists primarily of commercial banks, mutual funds and pension and insurance funds, holds a very large fraction of the outstanding mortgage bonds. It is very interesting that the fraction of bonds held by foreign investors is relatively low. It is under half of the share of government bonds held by foreign investors. The government bonds and the mortgage bonds share the foreign exchange risk. Hence, this cannot cause the large difference. The credit risk is of course different for mortgage bonds and government bonds, and this may in part cause the shares to diverge, if the foreign investors are relatively risk-averse. On the other hand, even though the Danish government debt is triple-A rated with both Moody's and Standard & Poor's, most of the newly issued Danish mortgage bonds also have a triple-A rating.[3] So, the credit risk can only to a very limited degree account for the difference in the fraction of bonds held by foreigners between government bonds and mortgage bonds.

[3] Danmarks Nationalbank (2005), p. 153 and Moody's (2005).

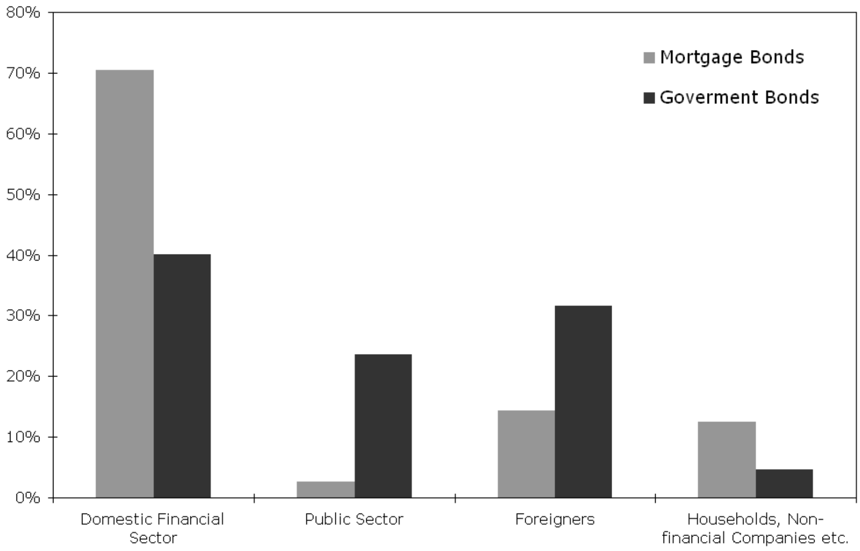

Source: Danmarks Nationalbank

Figure 1.2: Danish Bonds – investor distributions as of February 2006

Finally, the difference could also be caused by a difference in liquidity. However, most Danish government bonds and mortgage bonds are very liquid, and therefore we must expect the effect from this factor to be limited. Thus, this difference can primarily be attributed to the complexity of the construction of Danish mortgage bonds.[4]

1.2 Aim, Contribution and Literature

We present in this thesis, the important components of a pricing model for Danish mortgage bonds. The aim of the thesis is as follows:

> *To go through the various elements of pricing mortgage bonds with embedded prepayment options, namely we wish to develop and apply a term structure model and to discuss and set up a prepayment model. Furthermore, we seek to explain investment measures and strategies, dealing with mortgage bonds.*

The purpose of this thesis is to make a coherent presentation of how to value Danish mortgage bonds. We find that the literature on valuation of mortgage-backed products is vast, but very segmented. For instance, it is rare that both term structure modelling and prepayment modelling is treated in the same text. This is what we aim to do. Furthermore we strive at presenting it with a balance between academic and more practical

[4]Actually, a considerable share of the Danish mortgage bonds outstanding are short-term non-callable bonds, which are much more easy to price than callable bonds. Their existence is due to the introduction of adjustable rate mortgages on the Danish market. We return to this issue in section 9.

perspectives. Hence, the thesis should both be able to serve as an academic text on the valuation of mortgage bonds and as a more practically oriented text, which can be used as inspiration to the practical implementation of both term structure modelling and prepayment modelling. We try to give a comprehensive overview of the issues involved in mortgage bond valuation. Therefore, we also present a few sections on investment issues and a section on the product innovation on the Danish mortgage bond market to facilitate a more complete understanding of the mortgage bond pricing universe.

Hence, the combination of academic and practical aspects of mortgage bond pricing provides our contribution to the existing literature on mortgage bond pricing, which is after all relatively limited in the Danish context. Furthermore, we develop a new prepayment model, and in this connection we investigate whether preliminary prepayment data can provide an additional source of information when modelling prepayments. This approach is to our knowledge very briefly discussed in the existing literature. Furthermore, the development of investment strategies provides another contribution that distinguishes our thesis from much of the existing literature. This is due to the aforementioned symbiosis between academic and practical texts, which causes relevant practical issues to be illuminated in this thesis. This includes important considerations of choosing samples for estimating yield curves, obtaining prices and deciding on relevant calibrating instruments for the term structure model, deciding on the use of refinancing alternatives and interest rates etc.

The existing literature on the topic of pricing mortgage-backed products is mainly treating the American case. The Danish and American mortgage financing markets have many similarities and are globally unique due to the inclusion of a prepayment option.[5] Therefore, American literature on the topic can be used to a fairly high extent.[6] However, even though the Danish and the US markets share many features, there are still a few important differences. These will have little importance for the first part of the thesis, the modelling of the term structure, but are essential knowledge when treating the issue of prepayments. Therefore we save the presentation of the differences until the prepayment sections. The Danish literature on mortgage financing is relatively scarce, since much of the research conducted in this field is performed by quantitative units in commercial banks or specialized smaller companies, who have little interest in sharing their knowledge

[5]Other important markets are Netherlands, UK, Sweden and Germany, but the mortgage products are much more plain vanilla in these countries. See Miles (2003) for a nice comparison of the mortgage financing structures in countries in the European Union and the US.

[6]For an exhaustive text covering most issued involved in valuation of American mortgage-backed securities, consult Fabozzi (2001).

with others.

1.3 Structure

We present the structure of the thesis at hand by going through the principle of our mortgage bond pricing model, which is shown in Figure 1.3.

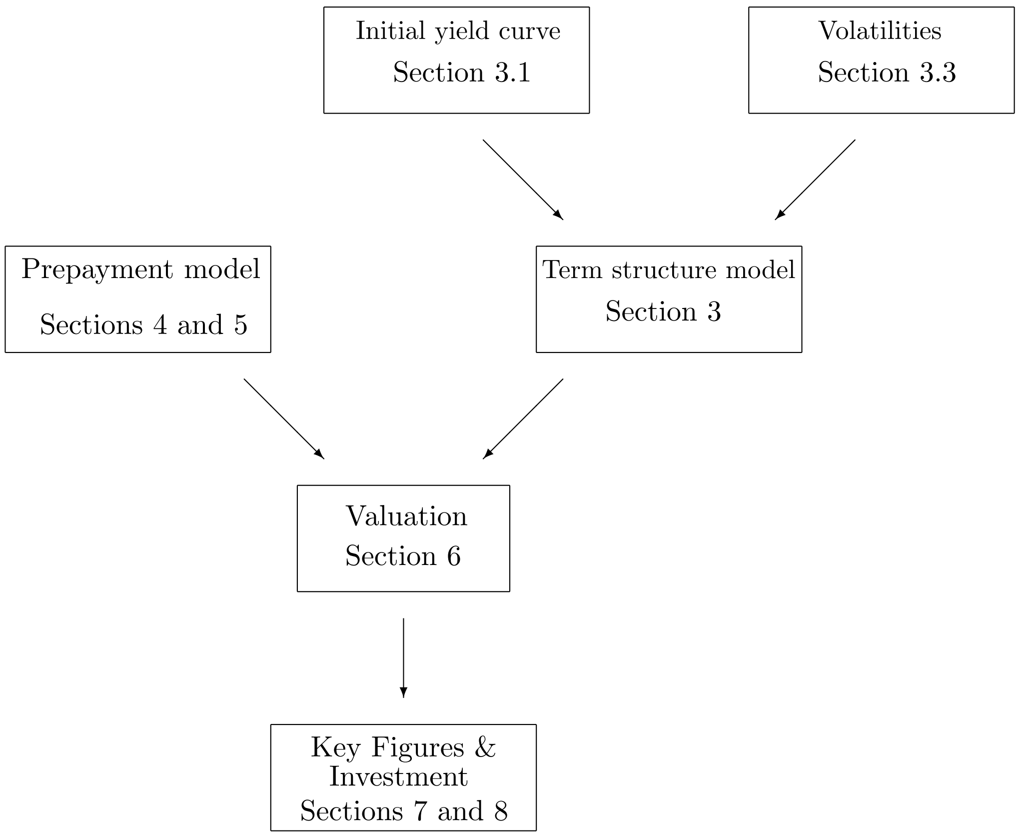

Figure 1.3: The structure of our mortgage bond pricing model.

As it is seen from the figure, the two main building blocks are a term structure model and a prepayment model. We start by treating the term structure model, afterwards turning to the modelling of prepayments.

However, before we start to set up and apply a term structure model, which we will do in section 3, we need to establish a mathematical pricing framework. This is the topic of section 2. This section may be skipped by the non-technically interested reader, even though we of course draw on the some of the results from this section in the subsequent sections. In section 3, we estimate an initial yield curve, we derive the pricing formulas of the chosen term structure model and we apply it to observed data. In section 4, we discuss prepayment behavior, leading to the modelling section 5, where we set up our own prepayment model. We close the pricing sections with section 6, where we outline the principle of combining the term structure model and the prepayment model. We subsequently turn towards the investment issue in sections 7 and 8, going through relevant return and risk measures and using these measures to create trading strategies. We close the thesis with a discussion of the product innovation of mortgage bond products that has been taking place during the last decade in section 9, and we conclude in section 10.

The issue of pricing mortgage bonds is, as mentioned earlier comprehensive. Therefore, we cannot treat every aspect involved in the process. We have chosen to be very thorough with the two main parts of the model, namely modelling the term structure and prepayments, while we have spent less effort on describing the process of combining these in practice. A thorough presentation of this would quickly turn into a question of technical and, in our opinion, less interesting issues. In this thesis, we focus specifically on the Danish mortgage bond market although we do make references to the US market in particular along the way. Furthermore, we only treat the issue of taxes to a very limited extent. Inclusion of taxes in the model would complicate things considerably, but actually bring few new qualitative insights, so we choose to discuss the issue of tax distortions only where we find it imperative.

2 Arbitrage-Free Pricing

In this section we construct the foundation for our term structure model. Due to the complexity of this field it has not been an integrated part of the finance courses available at the Department of Economics at University of Copenhagen. We therefore make an effort to present it such that a reader with the general background in financial economics finds it accessible.

The purpose of this section is to develop the general term structure equation, and furthermore to develop general formulas for bond and derivatives prices. We need these later in section 3, where we derive formulas for the chosen term structure model and use the derived derivatives prices to calibrate the term structure model. Furthermore, this section serves to facilitate the reader's understanding of arbitrage-free pricing and martingale probabilities in general.

We develop a pricing framework based on the assumption that markets are arbitrage-free following the line of thought that any arbitrage opportunities would be exploited instantly and hence be eliminated. The idea of arbitrage-free pricing is formalized by the assumption that if a market is arbitrage-free and a given claim Γ can be replicated, then the price of that claim at time t must be the replication cost $\Gamma(t)$. The concept of arbitrage free pricing can in general terms be written as

Definition 2.1 ***No Arbitrage Condition***
Any strategy with a zero initial investment cannot have a positive probability of producing a profit while at the same time having a zero probability of producing a loss.

We will also make the standard, but vital, assumption that markets are complete and frictionless. Completeness is a critical assumption. A market is complete if and only if every contingent claim is attainable[7] or in other words, any pay-off profile can be replicated using existing assets. The theory of arbitrage-free pricing hinges on claims being replicable. Not only is pricing in incomplete markets difficult to model, but incomplete markets can also change the model implications significantly. However, we maintain the assumption of complete markets, and refer to chapter 8 in Dana & Jeanblanc (2003) for an approach to cope with incomplete markets. To assume that markets are frictionless is somewhat innocent. More advanced models can easily take frictions into account, but in our case little is gained compared to the added complexity.

[7] Brigo & Mercurio (2001) p. 26.

2.1 Notation and Framework

Before we begin modelling the asset price, we introduce the probability theory framework. Probability theory is a vital element of asset pricing as we need to assess the likelihood of the occurrence of different states as this determines the value of the asset.

Say, we have a random variable X. For this variable we have a probability space denoted Ω, which is a set containing every value that X can take. A realized event is denoted ω where, of course, $\omega \in \Omega$. To this probability space belongs a probability measure, which we denote π.[8] The π-measure is called the objective or real-world measure. Say, X is the outcome of a single (fair) coin flip. In this example we have $\Omega = \{\text{Head}, \text{Tail}\}$, ω is either $\{\text{Head}\}$ or $\{\text{Tail}\}$ and $(\pi_H, \pi_T) = (\frac{1}{2}, \frac{1}{2})$.

To be theoretically stringent, we also use the notion of a sigma algebra $\Im(t)$ and a filtration $\{\Im(t)\}$, which are both functions of time, t. A sigma-algebra can be considered as being a set containing the information revealed by X. A filtration is then an increasing family of sigma-algebras, that is $\{\Im(0)\} \subseteq ... \subseteq \{\Im(t)\}$. To this we add the notion of a process being adapted. A process X is adapted if $X(t)$ is $\Im(t)$-measurable. This essentially means that given $\Im(t)$ we know $X(t)$ for any given t. We assume that the relevant processes are adapted, that is to assume complete information. To illustrate the concept of sigma algebras and filtrations, we now consider X to be an asset price. At time t we then have a sigma algebra $\Im(t) = \{X(t)\}$ which is a set containing the asset price at time t and the filtration $\{\Im(t)\} = \{X(s); 0 \leq s \leq t\}$ which is the entire price history up until time t. The four elements $(\Omega, \Im, \{\Im(t)\}, \pi)$ constitute together what is called a filtered space and it is within this framework that we construct our pricing model. This is a somewhat simple introduction to a very complicated field and we refer to Williams (1991) for a more thorough introduction to filtered spaces. We now move on to the development of our model, introducing the risk-free asset.

2.2 Money-Market Account

Let us start out by introducing the money-market account, which is the notion of the locally risk-free asset. It is a very simple, but also a very important asset as it allows us to relate cash flows across different points in time. One currency unit[9] invested, at any time t in the money-market account earns the prevailing instantaneous risk-free rate,

[8] In general, an unconstrained number of probability measures can be affiliated with a given probability space. See e.g. Williams (1991).

[9] Henceforth, we use \$ due to its notational convenience.

which can be written as[10]

$$dB(t) = B(t)r(t)dt \tag{2.1}$$

where $B(t)$ is the investment in the money-market account at time t. Hence, we see that the change in value of the money-market account over an infinitesimally small time interval, equals the deposit times the instantaneous risk-free interest rate. We assume for simplicity that $B(0) \equiv 1$, which implies that we obtain the discounting function as the inverse money-market account function. We have the following solution to the differential equation stated above

$$B(t, r(t)) = e^{\int_0^t r(s)ds} \tag{2.2}$$

We see that the value of the money-market account depends only on the evolution of the interest rate. The interest rate is generally the main source of variability in fixed income assets and we will give this area special attention in section 3. For now we merely assume that the interest rate can be represented under the objective measure by a general diffusion process also known as an Itô process[11]

$$dr(t) = \mu(t, r)dt + \sigma(t, r)dW \tag{2.3}$$

This process can parted into two terms; the $\mu(\cdot)$ function is called the drift term and the $\sigma(\cdot)$ function is called the diffusion term. The drift term indicates the deterministic part while the diffusion term indicates the stochastic part. The stochastic element originates from a Brownian motion denoted $W(t)$ and we state its definition below.

Definition 2.2 ***Brownian Motion***

A Brownian motion denoted W satisfies

- $W(0) = 0$
- *For any $0 \leq t_1 < t_2 < ... < t_n$ we have that $W(t_2) - W(t_1), ..., W(t_n) - W(t_{n-1})$ are independent*
- $W(t) - W(s) \sim N(0, \sqrt{t-s}), \quad \forall \quad t > s$
- *W has continuous sample paths*

[10] See e.g. Cairns (2004) p. 18.

[11] Denoting the interest rate process r is fairly unfortunate. In most other fields of economics r refer to the real interest rate and in most probability theory capital letters refer to variables while small letters refer to realized values of the process denoted by the capital letter. However, we follow the notation established in the literature.

The conditions ensure us that we have a random term that is continuous and satisfies the Martingale property. We return later in this section to what it means for a process to satisfy the Martingale property. Also notice that a Brownian motion is the limit case of a random walk. Though the process is continuous, it is not differentiable anywhere due to its irregularity. One can say that a Brownian motion over the next interval (no matter how small) can go anywhere as its increments are drawn from a normal distribution. These properties makes a Brownian motion highly suitable for modelling uncertainty.

Looking at (2.1), we can now see why the money-market account is only locally risk-free. Even though the interest rate is stochastic the investor knows the return he receives in the next instant (dt) with certainty. However, if an investor chooses to place a deposit in the money-market account over a discrete interval (Δt), the investment is no longer risk-free as the stochastic term in the interest rate process can carry the value of the investment in any direction. Thus, a deposit in the money-market account is considered to be only a locally risk-free asset. We use it as a numeraire asset when pricing risky assets.

2.3 Martingale Approach

There are several ways to derive a pricing model. We choose to apply the martingale approach as it is most commonly used. To understand why the martingale property is so applicable we start out by stating a very powerful theorem:[12]

Theorem 2.1 ***The First Fundamental Theorem of Asset Pricing***
A market is arbitrage-free if and only if there exists at least one equivalent martingale measure.[13]

Whilst the real-world uses a currency as numeraire, the martingale approach uses the risk-free asset. Essentially, the martingale approach assigns an equivalent probability measure to the probability space and applies this when pricing assets. Let us initially illustrate the (potential) difference between the objective measure and a pricing or market-implied measure by the following example.[14] Say, we have a two-period market with two assets and the interest rate equals zero for sake of simplicity. In the first period, we invest in an

[12]See e.g. Shreve (2004) p. 193.

[13]The equivalent martingale measure is sometimes also called the risk neutral measure or the risk adjusted measure. However, as the concept of martingales is the most precise of these names we stay with this name.

[14]Giesecke (2004) p. 33 provided inspiration for this example. An abridged version of Giesecke (2004) is published in Shimko (2004).

asset and in the second period we receive its pay-off. In the latter period, two states can occur; with probability $\pi_h = \frac{1}{2}$ a high return state and with $1 - \pi_h = \frac{1}{2}$ a low return state. As illustrated in Figure 2.1, asset A costs \$10 and pays \$10 in both states. Hence, asset A is risk-free. Asset B on the other hand costs \$5 and pays \$20 in state h and \$0 in state l.

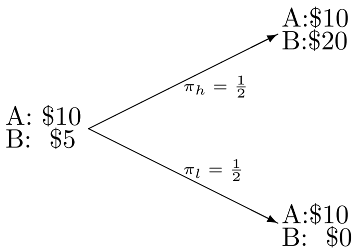

Figure 2.1: Asset pay-offs

From Figure 2.1, we can see that the objective measure is not used to price the assets as this would imply that asset B had an initial price of $\frac{1}{2}(20) + \frac{1}{2}(0) = 10$. Investors thus demand a risk premium to hold asset B equivalent to the difference between the time-0 price and the expected value. This implies a pricing probability of 25% for state h and 75% for state l as $0(Q_l) + 20(Q_h) = 5$ and $Q_h = 1 - Q_l$ provide us with $(Q_l, Q_h) = (\frac{3}{4}, \frac{1}{4})$. However, one should not confuse the Q-probabilities with the market belief of the likelihood of the two different states. They are merely a measure which incorporates a risk-adjustment for the uncertainty. That is, by pricing the asset using the Q-measure, which puts more weight on the low return state, we take into account that future returns are uncertain and that risk-averse investor demands a risk premium.

In this example we have not argued why the measure is an equivalent measure or why this measure is sensible to use, but we will do this as we are setting up the model and a more complete set of tools becomes available to us. Let us proceed by defining what is meant by the equivalent martingale measure.

Definition 2.3 ***Equivalent Measures***

Say we have two probability measures, π and Q, on $(\Omega, \Im)$. These two measures are said to be equivalent measures if for any event $\omega \in \Omega$

$$\pi(\omega) > 0 \quad \Leftrightarrow \quad Q(\omega) > 0 \tag{2.4}$$

The definition of an equivalent measure is thus very straightforward. In simple terms, it requires the measures to agree on which events have respectively positive and zero probability, but does not impose any restrictions on the relative likelihood of the events. But for an equivalent measure to also be a pricing measure it must additionally be a martingale measure. This is the case if the discounted asset prices are martingales under the equivalent measure and we, therefore, define what is required for a process to be a martingale process.

Definition 2.4 ***Martingale***
Given a probability triple $(\Omega, \Im, Q)$*, the adapted process* X *is called a martingale (relative to* $(\Im, Q)$*) if*

- $E^Q[|X(t)|] < \infty$
- $E^Q[X(t)|\Im(s), s \leq t] = X(s)$

The first condition is of technical matter while the second condition is the main property of a martingale. It states that given we know $X(s)$ at time s and we wish to estimate a future $X(t)$, then our best estimate is indeed $X(s)$. This is equivalent to the expected change in X being zero.

We now go through the set-up of an arbitrage-free pricing model using the equivalent martingale measure.[15] We assume that the bond is free of credit risk, which implies that a payment of a coupon at time t does not affect borrower's ability to pay future claims. We can, therefore, view a coupon bond as a portfolio of zero coupon bonds and hence it suffices to price zero coupon bonds.

Say we have a zero coupon bond paying \$1 at time T (a so-called T-bond) that we wish to price. We assume that the price have the following form

$$p(t, T, 1) = F(t, r(t); T) \tag{2.5}$$

with

$$F(T, r; T) = 1 \tag{2.6}$$

where F is continuous and of class C^1 with respect to t and C^2 with respect to r. We search for a general function which is only restricted by loose regularity conditions. Of course, assuming no credit risk, we also impose the condition that the bond surely pays \$1 at maturity regardless of which r is realized.

[15] This presentation of arbitrage free pricing is partly based on Björk (1998).

We now investigate which restrictions we must place on the dynamics of F in order for it to be an arbitrage-free price. We apply Itô's Lemma to the bond price, F, and get[16]

$$dF = \frac{\partial F}{\partial t}dt + \frac{\partial F}{\partial r}dr + \frac{1}{2}\frac{\partial^2 F}{\partial r^2}d\langle r\rangle \tag{2.7}$$

where, by use of the convention that $dWdW = dt$ and $dWdt = dt^{\frac{3}{2}}$,[17] we get

$$\begin{aligned} d\langle r\rangle &\equiv drdr \\ &= (\mu(t,r)dt + \sigma(t,r)dW)(\mu(t,r)dt + \sigma(t,r)dW) \\ &= \mu^2(t,r)(dt)^2 + \sigma^2(t,r)dt + 2\mu(t,r)\sigma(t,r)(dt)^{\frac{3}{2}} \end{aligned}$$

As we are working in a continuous framework we are letting dt become infinitesimally small. Therefore, we ignore terms containing $(dt)^2$ and $(dt)^{\frac{3}{2}}$ as they go to zero faster than dt. Thus, we have

$$d\langle r\rangle = \sigma^2(t,r)dt \tag{2.8}$$

We obtain the following expression for the change in the asset price by inserting (2.3) and (2.8) into (2.7)

$$\begin{aligned} dF &= \frac{\partial F}{\partial t}dt + \frac{\partial F}{\partial r}[\mu(t,r)dt + \sigma(t,r)dW)] + \frac{1}{2}\frac{\partial^2 F}{\partial r^2}\sigma^2(t,r)dt \\ &= \left[\frac{\partial F}{\partial t} + \mu(t,r)\frac{\partial F}{\partial r} + \frac{1}{2}\frac{\partial^2 F}{\partial r^2}\sigma^2(t,r)\right]dt + \sigma(t,r)\frac{\partial F}{\partial r}dW \end{aligned} \tag{2.9}$$

This can be rearranged into a geometric Brownian motion

$$dF = \tilde{\alpha}(t,r)Fdt + \tilde{\sigma}(t,r)FdW \tag{2.10}$$

where

$$\tilde{\alpha}(t,r) = \frac{\frac{\partial F}{\partial t} + \mu(t,r)\frac{\partial F}{\partial r} + \frac{1}{2}\sigma^2(t,r)\frac{\partial^2 F}{\partial r^2}}{F} \tag{2.11}$$

$$\tilde{\sigma}(t,r) = \frac{\sigma(t,r)\frac{\partial F}{\partial r}}{F} \tag{2.12}$$

[16] A word on the notation; to simplify the mathematical expressions throughout the derivation we omit function arguments whenever it seems fitting; e.g. if no principal is included in $p(t,T)$ it is the price of a $1 bond.

[17] Dana & Jeanblanc (2003) p. 86.

We have now calculated the dynamics of the bond price, F, and using these, we do the same for the discounted price process. Recall from the definition of the equivalent martingale measure that we require that the discounted price process is a martingale in order for us to apply the First Theorem of Asset Pricing. We derive the discounted price process as

$$Z(t,T) = F(t,T)B^{-1}(t) \tag{2.13}$$

Notice that by working with the discounted price process, we are changing numeraire from currency units to money-market account units. By applying the product rule for stochastic processes[18], we get

$$dZ(t,T) = dF(t,T)B^{-1}(t) + F(t,T)dB^{-1}(t) + d\langle F(t,T), B^{-1}(t)\rangle \tag{2.14}$$

To determine the dynamics of the discounted asset price we first need to determine

$$\begin{aligned} dB^{-1}(t) &= -B^{-2}(t)dB(t) + \frac{2}{B^3(t)}d\langle B(t)\rangle \\ &= -r(t)dtB^{-1}(t) + \frac{2}{B^3(t)}(B(t)r(t)dt)^2 \\ &= -r(t)dtB^{-1}(t) \end{aligned} \tag{2.15}$$

$$\begin{aligned} d\langle F(t,T), B^{-1}(t)\rangle &= dFdB^{-1}(t) \\ &= -[\tilde{\alpha}(t,r)Fdt + \tilde{\sigma}(t,r)FdW]r(t)dtB^{-1}(t) \\ &= 0 \end{aligned} \tag{2.16}$$

Both results stem from the previously mentioned fact that dt of a higher order than 1 can be ignored. We can then rewrite (2.14) by inserting (2.15) and (2.16) as

$$\begin{aligned} dZ(t,T) &= dF(t,T)B^{-1}(t) - F(t,T)r(t)dtB^{-1}(t) \\ &= F(t,T)B^{-1}(t)[\tilde{\alpha}(t,r)dt + \tilde{\sigma}(t,r)dW] - F(t,T)B^{-1}(t)r(t)dt \\ &= Z(t,T)[(\tilde{\alpha}(t,r) - r(t))dt + \tilde{\sigma}(t,r)dW] \end{aligned} \tag{2.17}$$

We have now calculated the dynamics of the discounted price process and are interested in finding the equivalent measure under which it is a martingale; that is where $dZ(t,T)$

[18] If $Z = X(t)Y(t)$ then $dZ = XdY + YdX + d\langle X, Y\rangle$.

has a zero drift term. We therefore define

$$\lambda(t) \equiv \frac{\tilde{\alpha}(t,r) - r(t)}{\tilde{\sigma}(t,r)} \tag{2.18}$$

which is recognized as the market price of risk. It is the excess return over risk-free rate per unit of risk received by investor for holding the (interest rate) risky T-bond. Notice that if investors are risk averse then $\lambda(\cdot)$ is positive as people will demand a risk premium to hold a risky asset. Assuming $\lambda(\cdot)$ satisfies the technical Novikov condition, we can apply the Girsanov Theorem, which ensures the existence of the Q-measure.[19] The theorem also provides us with the Brownian motion under the Q-measure.

$$\begin{aligned} W^Q(t) &\equiv W(t) + \int_0^t \lambda(s)ds \Rightarrow \\ dW^Q(t) &= dW(t) + \lambda(t)dt \end{aligned} \tag{2.19}$$

To verify that the Q-measure is the martingale measure, we insert (2.19) into (2.17)

$$\begin{aligned} dZ(t,T) &= Z(t,T)[\tilde{\alpha}(t,r) - r(t) - \lambda(t)\tilde{\sigma}(t,r)dt] + \tilde{\sigma}(t,r)(dW + \lambda(t)dt) \\ &= Z(t,T)\tilde{\sigma}(t,r)dW^Q \end{aligned} \tag{2.20}$$

Note that (2.20) only includes a diffusion term of which the expected value under Q equals zero. $Z(\cdot)$ is then a martingale subject to the technical condition from Definition 2.4, that is $E^Q[\exp(\frac{1}{2}\int_0^T \tilde{\sigma}^2(t,r)dt)] < \infty$.

Björk (1998) shows that if markets are arbitrage-free, then there exists a stochastic process $\lambda(t)$ as defined above for each maturity T. Björk (1998), furthermore, shows by constructing a perfectly hedged portfolio consisting of two risky bonds with maturities T and S that $\lambda_T(\cdot) = \lambda_S(\cdot)$. It makes intuitive sense that any two bonds, regardless of their maturities, necessarily have the same market price of risk or equivalently risk-adjusted return if the market is arbitrage-free.

By now we know that $\lambda(\cdot)$ is universal and must therefore be identical for all maturities. Therefore, we can apply (2.18) to derive the term structure equation which is the a no-arbitrage condition for the dynamics of our asset price. Inserting (2.11) and (2.12) into

[19] Cairns (2004) p. 248.

(2.18) gives us

$$\begin{aligned}
\frac{\frac{\frac{\partial F}{\partial t}+\mu(t,r)\frac{\partial F}{\partial r}+\frac{1}{2}\sigma^2(t,r)\frac{\partial^2 F}{\partial r^2}}{F}-r(t)}{\frac{\sigma(t,r)\frac{\partial F}{\partial r}}{F}} &= \lambda(t) \Rightarrow \\
\frac{\partial F}{\partial t}+\mu(t,r)\frac{\partial F}{\partial r}+\frac{1}{2}\sigma^2(t,r)\frac{\partial^2 F}{\partial r^2}-r(t)F &= \lambda(t)\sigma(t,r)F_r \Rightarrow \\
\frac{\partial F}{\partial t}+[\mu(t,r)-\lambda(t)\sigma(t,r)]\frac{\partial F}{\partial r}+\frac{1}{2}\sigma^2(t,r)\frac{\partial^2 F}{\partial r^2}-r(t)F &= 0
\end{aligned}$$

We also need to include the boundary condition stated in (2.6). Together these two equations provide us with the important result.

Result 2.1 ***Term Structure Equation***

In an arbitrage-free market a zero coupon bond price denoted $F(t,T)$ must satisfy

$$\frac{\partial F}{\partial t}+[\mu(t,r)-\lambda(t)\sigma(t,r)]\frac{\partial F}{\partial r}+\frac{1}{2}\sigma(t,r)^2\frac{\partial^2 F}{\partial r^2}-r(t)F = 0 \quad (2.21)$$

$$F(T;T) = 1 \quad (2.22)$$

Notice the resemblance to the famous Black-Scholes partial differential equation.[20] However, the term structure equation is more complex due to the occurrence of $\lambda(\cdot)$ (and a stochastic interest rate). As an aside, we briefly discuss this added complexity as it has been important for the modelling of term structure models.[21] To see how the objective measure relates to the equivalent martingale measure we use the Radon-Nikodym density

$$\begin{aligned}
V &= \frac{dQ}{d\pi}|_{\Im(t)} \\
&= e^{-\frac{1}{2}\int_0^t \lambda^2(s)ds-\int_0^t \lambda(s)dW}
\end{aligned} \quad (2.23)$$

As it can be seen from (2.23), the density is essentially a likelihood ratio. Furthermore, by setting $\lambda(\cdot)=0$ in (2.23), which is equivalent to investors being risk-neutral, we see that $V=1$, hence, no adjustment is needed. In other words, the equivalent martingale measure is indeed the risk-neutral measure.

We now apply the Radon-Nikodym density to the introductory example to show how this random variable describes the relationship between π and Q. In the discrete case the

[20] See e.g. Cvitanic & Zapatero (2004) p. 223.

[21] It would also be important for the Black-Scholes-Merton framework had it not included additionally simplifying assumptions such as a constant risk-free interest rate.

density function is point probabilities which gives us

$$V(\omega_h) = \frac{.25}{.5} = .5 \quad \text{and} \quad V(\omega_l) = \frac{.75}{.5} = 1.5$$

Hence, given risk-aversion (implied by the time-0 asset prices) the Q-measure mitigates the expected growth rates as we argued in the introductory example. This can also be seen looking at the interest rate and price process under the equivalent martingale measure. The interest rate evolves according to

$$dr(t) = [\mu(t,r) - \lambda(t)\sigma(t,r)]dt + \sigma(t,r)dW^Q \tag{2.24}$$

and the price evolves according to

$$dF(t) = F(t,T)[r(t)dt + \tilde{\sigma}(t,r)dW^Q] \tag{2.25}$$

where the asset price now has a drift rate equal to the risk-free rate. Note also that by changing the numeraire we are only affecting the drift term. Hence, when changing measure one does not need to change ones volatility process.

As $\lambda(\cdot)$ is not determined endogenously we would need to define it exogenously. We can see that our choice of $\lambda(\cdot)$ dictates how we move from the objective measure to the Q-measure or vice versa. One way of avoiding the troubles of modelling $\lambda(\cdot)$ (explicitly) is to model the interest rate process directly under Q. When pricing interest rate derivatives we do not need to move from the Q-measure to the objective measure as market prices are observed as expectations under the Q-measure. As explained in the introductory example, assets are not priced according to the objective measure, but instead under the pricing measure Q. Indeed, as we will see in Section 3, it is standard practice to model the interest rate process directly under Q.[22]

We have now established the existence of an equivalent martingale measure under which we can price financial assets. We have thus obtained an important result for our aspirations of pricing mortgage bonds.

[22] We also have to calibrate the model using Q-dynamics. We will carry out the calibration in the next section.

Result 2.2 ***General Pricing Formula***

For each time t, $0 \leq t \leq T$*, there exists a unique price*[23]

$$F(t,T,\Gamma) = E^Q[e^{-\int_t^T r(s)ds}\Gamma|\Im(t)] \tag{2.26}$$

for a given attainable claim of $\$\Gamma$ *with maturity* T*.*[24]

We can then easily obtain the price of a zero coupon bond with a principal L by replacing Γ with L. Not surprisingly, the price of such a bond is

$$\begin{aligned} p(t,T,L) &= E^Q[e^{-\int_t^T r(s)ds}L|\Im(t)] \\ &= L \cdot p(t,T) \end{aligned} \tag{2.27}$$

Though our main pricing formula is easily applicable to simple claims, we run into trouble when Γ and $r(t)$ are dependent. Notice that even if Γ and $r(t)$ are independent under the objective measure, they are still dependent under Q. It can be seen from (2.25) that any asset has a local drift rate under Q equal to the risk-free rate. The well-known framework of Black-Scholes-Merton assumes a constant risk-free rate and hence avoids the problem of dependence between the two variables. Beyond such simplifications, we would need to know the joint distribution of the two variables under Q in order to calculate the expectation in our general pricing function.

We are, therefore, going to extend our general pricing formula such that it can be applied to more advanced forms of pay-off profiles. We would like to be able to write (suppressing the conditioning)

$$\begin{aligned} F(t,T,\Gamma) &= E[e^{-\int_t^T r(s)ds}]E[\Gamma] \\ &= p(t,T)E[\Gamma] \end{aligned} \tag{2.28}$$

This would obviously be desirable, as the expectation we need to calculate becomes relatively simple and we can, at time t, observe $p(t,T)$. To facilitate this, we define a new probability measure[25]

[23] Uniqueness of price relies on the assumption that the claim is attainable. From the second fundamental theorem of pricing we know that the equivalent martingale measure Q is unique if and only if markets are complete.

[24] Brigo & Mercurio (2001) p. 26.

[25] We refer to Björk (1998) p. 283-285 for a proof of the existence of the measure.

Definition 2.5 ***Forward Measure***

Say we have an arbitrary T-maturity claim Γ *and a T-bond with price* $p(t,T)$*. The T-forward neutral measure,* Q^T*, then allows us to write*

$$\Upsilon(t;\Gamma) = p(t,T)E^T[\Gamma] \tag{2.29}$$

where Υ *denotes the forward measure value*[26] *and* $E^T[\cdot]$ *denotes the expectation under the T-measure. The T-bond is called the numeraire of the forward measure.*

To show how the additional change of numeraire applies, we go through the pricing of more advanced assets – European call and put options on a zero coupon bond. Conveniently, we need these formulas later on in section 3.3. The two options share some general characteristics such as the underlying asset is a zero coupon bond paying \$$L$ at time T_2. The options have maturity T_1 (where of course $T_1 \leq T_2$) and a strike price of K. We can thus write the pay-off of the call option as

$$\max[0, p(T_1,T_2,L) - K] = [L \cdot p(T_1,T_2) - K]^+ \tag{2.30}$$

and the pay-off for the put option as

$$\max[0, K - p(T_1,T_2,L)] = [K - L \cdot p(T_1,T_2)]^+ \tag{2.31}$$

As the derivation of prices for the two options are very similar, we are only going through the technique for a call option and we merely state the price of the put option. In order to value the call option on the T_2-bond, we apply (2.26) to find the option price ZBC

$$\begin{aligned} ZBC(t,T_1,T_2,K,L) &= E^Q\left[e^{-\int_t^{T_1} r(s)ds}[L \cdot p(T_1,T_2) - K]^+ | \Im(t)\right] \\ &= E^Q\left[e^{-\int_t^{T_1} r(s)ds} L \cdot p(T_1,T_2)\mathbf{1}_{\{L\cdot p(T_1,T_2)>K\}} | \Im(t)\right] \\ &\quad - E^Q\left[e^{-\int_t^{T_1} r(s)ds} K \cdot \mathbf{1}_{\{L\cdot p(T_1,T_2)>K\}} | \Im(t)\right] \end{aligned} \tag{2.32}$$

where the indicator function, $\mathbf{1}_\omega$, is equal to one if event ω occurs and zero otherwise. By further rearranging and changing numeraire from the Q-measure to the relevant maturity

[26] $\Upsilon(\cdot)$ is the T-measure equivalent to $F(\cdot)$ under the Q-measure.

forward measures we obtain

$$\begin{aligned} ZBC(t,T_1,T_2,K,L) &= E^Q\left[L \cdot e^{-\int_t^{T_2} r(s)ds}\mathbf{1}_{\{L \cdot p(T_1,T_2)>K\}}|\Im(t)\right] - \\ &\quad K \cdot E^Q\left[e^{-\int_t^{T_1} r(s)ds}\mathbf{1}_{\{Lp(T_1,T_2)>K\}}|\Im(t)\right] \\ &= L \cdot p(t,T_2)Q^{T_2}\left\{L \cdot p(T_1,T_2) > K\right\} - \\ &\quad K \cdot p(t,T_1)Q^{T_1}\left\{L \cdot p(T_1,T_2) > K\right\} \end{aligned} \tag{2.33}$$

It is now a matter of calculating the forward neutral probabilities in order to price the call option. The standard condition for this to be possible is that the volatility term is deterministic. We refer to Björk (1998) for the proof.

We start out by deriving the latter of the two probabilities in (2.33). As we are working under the T_1-measure (with the T_1-bond as a numeraire) let us define

$$M(t) \equiv \frac{p(t,T_2)}{p(t,T_1)} \tag{2.34}$$

which we assume evolves according to

$$dM(t) = M(t)[m(t)dt + \sigma_M(t)dW] \tag{2.35}$$

Furthermore, we assume that $\sigma_M(t)$ is deterministic such that we obtain computability, but we will need to check this assumption when using a particular price process later on. We conveniently use (2.34) to redefine the probability under the T_1-measure as

$$\begin{aligned} Q^{T_1}\left\{L \cdot p(T_1,T_2) > K\right\} &= Q^{T_1}\left\{\frac{L \cdot p(T_1,T_2)}{p(T_1,T_1)} > K\right\} \\ &= Q^{T_1}\left\{L \cdot M(T_1) > K\right\} \end{aligned} \tag{2.36}$$

We have now redefined the probability, such that it depends on the distribution of M under the forward neutral measure. To find the distribution of M, we must start out by deriving the dynamics of M under the T_1-measure.

$M(t)$ is an asset price normalized by the T_1-bond and thus, it has zero drift under Q^{T_1} and we can write its Q^{T_1}-dynamics as

$$dM(t) = M(t)\sigma_M(t)dW^{T_1} \tag{2.37}$$

This is a geometric Brownian motion and the solution to the differential equation is

$$\begin{aligned} M(T_1) &= M(t)e^{-\frac{1}{2}\int_t^{T_1}\sigma_M^2(s)ds+\int_t^{T_1}\sigma_M(s)dW^{T_1}} \\ &= M(t)e^{\kappa} \quad , \quad \kappa \equiv -\frac{1}{2}\int_t^{T_1}\sigma_M^2(s)ds + \int_t^{T_1}\sigma_M(s)dW^{T_1} \end{aligned} \tag{2.38}$$

We can see that it is the exponent κ that determines the distribution of M. We note that it contains two terms, which are respectively a deterministic integral and a stochastic integral. As the stochastic integral is a continuous summation of Brownian motion increments with a deterministic coefficient, it has the following distribution[27]

$$\int_t^{T_1}\sigma_M(s)dW_1^T \sim N\left(0,\Sigma^2\right), \quad \Sigma^2 \equiv \int_t^{T_1}\sigma_M^2(s)ds \tag{2.39}$$

We then correct the mean by the deterministic integral and subsequently normalize the variable by which we obtain[28]

$$\begin{aligned} \kappa &\sim N\left(-\frac{1}{2}\underbrace{\int_t^{T_1}\sigma_M^2(s)ds}_{=\Sigma^2} \quad , \quad \Sigma^2\right) \Longrightarrow \\ \frac{\kappa-\frac{1}{2}\Sigma^2}{\sqrt{\Sigma^2}} &\sim \Phi \end{aligned}$$

where Φ denotes the standardized normal distribution. We have now obtained the distribution of κ, which enables us to calculate the probability in (2.36) as

$$\begin{aligned} Q^{T_1}\left\{L\cdot M(T_1) > K\right\} &= Q^{T_1}\left\{\frac{L\cdot p(t,T_2)}{p(t,T_1)}e^{\kappa} > K\right\} \\ &= Q^{T_1}\left\{\kappa < \ln\left(\frac{L\cdot p(t,T_2)}{K\cdot p(t,T_1)}\right)\right\} \\ &= \Phi(h-\Sigma) \end{aligned} \tag{2.40}$$

where

$$h = \frac{1}{\Sigma}\ln\left[\frac{L\cdot p(t,T_2)}{K\cdot p(t,T_1)}\right] + \frac{\Sigma}{2} \tag{2.41}$$

In a similar fashion, we find

$$Q^{T_2}\left\{L\cdot M(T_1) > K\right\} = \Phi(h) \tag{2.42}$$

[27]Björk (1998), p. 43.
[28]Ruppert (2004) p. 15.

We can now calculate the option price by inserting (2.40) and (2.42) into (2.33).

Result 2.3 ***Zero Coupon Bond Option Prices***

The price of an European call option on a zero-coupon bond paying \$L at time T_2 where the option has maturity T_1 and a strike price of K can be written as

$$ZBC(t, T_1, L) = L \cdot p(t, T_2)\Phi(h) - K \cdot p(t, T_1)\Phi(h - \Sigma) \tag{2.43}$$

The price of the European put option on the same bond with the same strike price and maturity can be written as

$$ZBP(t, T_1, L) = K \cdot p(t, T_1)\Phi(-h + \Sigma) - L \cdot p(t, T_2)\Phi(-h) \tag{2.44}$$

where h is defined as in (2.41)

In this section, we have established the foundation for pricing of an attainable claim. Using the general model we have demonstrated how to use the technique to price simple as well as more complex claims. Most importantly, we used the martingale approach to derive the term structure equation, which provides us with the dynamics of the asset price for a given term structure model. In the next section, we model the term structure of interest rates, where we make use for the term structure equation as well as the derived expressions for prices of options on zero coupon bonds.

3 Term Structure Model of Interest Rates

Now that we have completed the needed general pricing set-up, we turn our view towards the term structure of interest rates. We proceed as follows: In section 3.1, we start by discussing how to actually obtain an initial term structure, before we turn to the issue of how to model the future evolution of the term structure in section 3.2. In that section we discuss various possible models of the term structure and derive the pricing formulas in the chosen model. We proceed to calibrate the parameters of the model in section 3.3, before we show how to apply the term structure model in section 3.4, using the estimated initial term structure and the calibrated model parameters.

3.1 Initial Yield Curve

This subsection deals with the issue of how to derive an initial term structure (yield curve). We need it later when applying the term structure model in section 3.4.

3.1.1 Yield Curve Modelling

The issue of estimating an initial yield curve has long been an issue in finance theory that has received much attention. Before one starts to address the issue of how to estimate a yield curve, one has to be sure exactly what is meant by this. Normally, when referring to the yield of a bond, one is talking about the yield to maturity, namely the discount rate that makes the present value of a payment stream equal to the price of the bond[29]

$$P = \sum_{t=1}^{T} \frac{CF_t}{(1 + \frac{\text{yield}}{\text{frq}})^{t \cdot \text{frq}}} \tag{3.1}$$

It would be tempting to draw the yields to maturity of various bonds in a maturity-yield space, and estimate a term structure on basis of this. However, that would be very misleading. In the calculation of the yield to maturity, all payments are assumed to be discounted with the same interest rate. The yield to maturity is therefore by definition constant over the lifetime of a bond. Hence, it can be regarded as some kind of an average interest rate of all the coupon payments made along the maturity of the bond. Alas, the yield to maturity is generally regarded as a unsatisfactory measure of the term structure, and often other interest rates than the yield to maturity are used to characterize the term structure. The most commonly used interest rates when describing the term structure

[29]Grinblatt & Titman (2002), p. 58.

are the zero coupon interest rates. The differences between the yield to maturity and the zero coupon interest rates are summarized below:

– The yield to maturity is the same for all payments of a bond – regardless of the timing of the various payments, but is typically different for different bonds, while

– The zero coupon interest rate is the same for all payments that mature at a specific point in time – regardless of which bond is under consideration, but is typically different for different maturities.[30]

Note that the yield to maturity and the zero coupon interest rate curves coincide in the particular case of a completely flat term structure. If the zero coupon interest rates are constant for all maturities, this constant value must exactly equal any weighted average of these interest rates; thus also the yield to maturity.

Sometimes, the forward rates ($f(m)$) are used in lieu of the zero coupon interest rates ($r(m)$). Fortunately, these two interest rates are closely related,

$$r(m) = \frac{1}{m}\int_0^m f(x)dx \tag{3.2}$$

such that it is easy to calculate one of them once you know the other.

To estimate a term structure of zero coupon interest rates, we need a model. Nelson & Siegel (1987) note that Durand (1942) was one of the first to make a suggestion in this direction. His suggestion was to "*draw a monotonic envelope under the scatter of points in a way that seemed to him subjectively reasonable*".[31] Since then, many researchers have tried to come up with better explanations and models for the initial term structure. For many years now, it has been widely accepted that the models used to estimate the initial term structure are simply statistically motivated, and usually with little economic content. Estimating the current term structure is simply a matter of reaching a parsimonious model that fits data well.

One of the most simple ways to combine the points into a term structure is the method of bootstrapping. This method applies linear interpolation to obtain a fully specified yield curve. This is usually regarded as a too simplistic method to estimate a yield curve. On the other hand, bootstrapping has an advantage in the fact that it per se fits data perfectly.

Another class of models often used in practice have been the so-called *spline* models, for instance the *cubic spline* model. The idea in these models is to divide the maturity span

[30] Christensen (2005), p. 47.
[31] Nelson & Siegel (1987), p. 474.

into smaller segments, and fit data to segmented polynomial curves, which are *splined* in fixed knots. This allows for a high degree of flexibility when estimating the term structure, but the procedure requires quite a lot of observations in order to make a robust estimation, especially if the number of knots is large.

We choose to apply a different class of models, namely the class of models originating in Nelson & Siegel (1987). This choice is supported by the findings in BIS (1999), which investigates the use of term structure models in a selection of central banks.[32] The result is that almost all of the central banks use the Nelson & Siegel (1987) model or the Svensson (1994) extension thereof. Only two of the central banks used a smoothing spline approach. Hence, this motivates us to proceed with the Nelson-Siegel and Svensson models.

At first, the Nelson-Siegel model assumes that the instantaneous forward rate at maturity m, which is denoted $f(m)$, is given by the solution to a second-order differential equation, which has two different real roots:

$$f(m) = \beta_0 + \beta_1 \cdot e^{-\frac{m}{\tau_{1a}}} + \beta_2 \cdot e^{-\frac{m}{\tau_{1b}}} \tag{3.3}$$

They investigate this model, and they find that this model is over-parameterized in their samples. This leads them to suggest a more parsimonious model with equal roots given by:

$$f(m) = \beta_0 + \beta_1 \cdot e^{-\frac{m}{\tau_1}} + \beta_2 \cdot \frac{m}{\tau_1} \cdot e^{-\frac{m}{\tau_1}} \tag{3.4}$$

Hence, from (3.3) to (3.4), the number of parameters is reduced from five to four. This model can, even though it has a relatively small number of parameters, generate various different shapes of term structures, including humps, S shapes, and monotonic curves, and provides often a reasonably good fit.[33] Various possible shapes of the yield curve under the Nelson-Siegel functional form are illustrated in Figure 3.1.

Svensson (1994) proposed an extension of the Nelson-Siegel model, which is very much used, by adding one term with two new parameters to (3.4):

$$f(m) = \beta_0 + \beta_1 \cdot e^{-\frac{m}{\tau_1}} + \beta_2 \cdot \frac{m}{\tau_1} \cdot e^{-\frac{m}{\tau_1}} + \beta_3 \cdot \frac{m}{\tau_2} \cdot e^{-\frac{m}{\tau_2}} \tag{3.5}$$

The purpose of making this extension was primarily that it allowed for a second *hump* in the term structure. Furthermore, Svensson (1994) found that it improved the fit in *his* sample substantially. In general, which of these models one prefers, is the standard trade

[32]The sample consists of a range of European countries, Japan and United States.

[33]Nelson & Siegel (1987), p. 476.

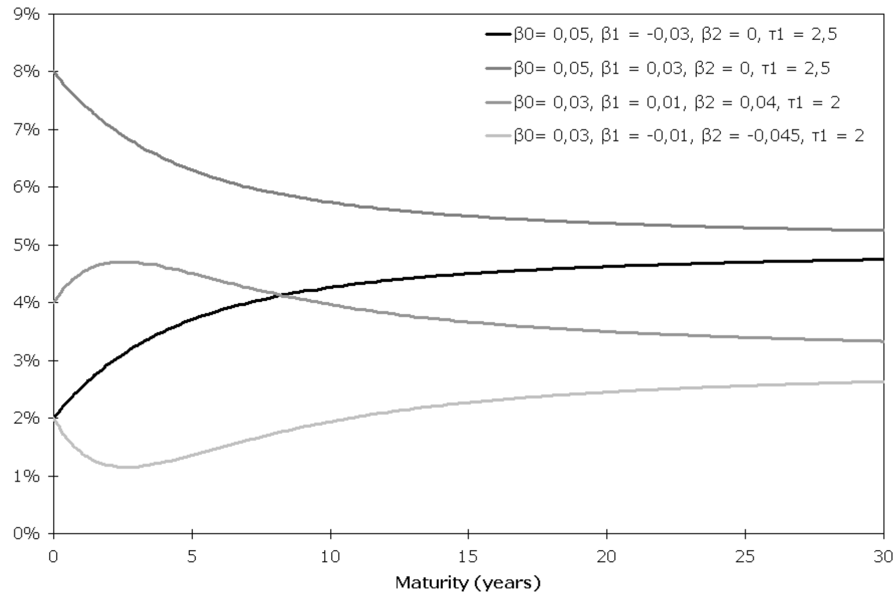

Figure 3.1: Various possible shapes of the yield curve in the Nelson-Siegel model

off between an improved fit and a parsimonious model, and it must be determined from case to case.

We now transform the model for the instantaneous forward rates into a model for the yield as measured by the zero coupon interest rates (the term structure). This is done by integrating the equation for the instantaneous forward rate from 0 to m and dividing by m as in (3.2). Since (3.3) and (3.4) are merely special cases of (3.5), we show how to derive an equation for the zero coupon interest rates based on (3.5).

$$\begin{aligned}
r(m) &= \left(\int_0^m \left[\beta_0 + \beta_1 \cdot e^{-\frac{x}{\tau_1}} + \beta_2 \cdot \frac{x}{\tau_1} \cdot e^{-\frac{x}{\tau_1}} + \beta_3 \cdot \frac{x}{\tau_2} \cdot e^{-\frac{x}{\tau_2}}\right] dx\right) / m \quad \Leftrightarrow \\
&= \left(\beta_0 \cdot m + \beta_1 \cdot \int_0^m e^{-\frac{x}{\tau_1}} dx + \frac{\beta_2}{\tau_1} \cdot \int_0^m x \cdot e^{-\frac{x}{\tau_1}} dx + \frac{\beta_3}{\tau_2} \cdot \int_0^m x \cdot e^{-\frac{x}{\tau_2}}\right) / m \quad \Leftrightarrow \\
&= \left(\beta_0 \cdot m + \beta_1 [-\tau_1 \cdot e^{-\frac{x}{\tau_1}}]_0^m + \frac{\beta_2}{\tau_1}\left([-\tau_1 \cdot e^{-\frac{x}{\tau_1}} \cdot x]_0^m - \int_0^m (-\tau_1 \cdot e^{-\frac{x}{\tau_1}}) dx\right)\right. \\
&\quad \left. + \frac{\beta_3}{\tau_2}\left([-\tau_2 \cdot e^{-\frac{x}{\tau_2}} \cdot x]_0^m - \int_0^m (-\tau_2 \cdot e^{-\frac{x}{\tau_2}}) dx\right)\right) / m \quad \Leftrightarrow \\
&= \Big(\beta_0 \cdot m - \beta_1 \tau_1 e^{-\frac{m}{\tau_1}} + \beta_1 \tau_1 - \beta_2 m e^{-\frac{m}{\tau_1}} - \beta_2 \tau_1 e^{-\frac{m}{\tau_1}} + \beta_2 \tau_1 \\
&\quad - \beta_3 m e^{-\frac{m}{\tau_2}} - \beta_3 \tau_2 e^{-\frac{m}{\tau_2}} + \beta_3 \tau_2\Big) / m \quad \Leftrightarrow \\
&= \beta_0 + \beta_1 \frac{1 - e^{-\frac{m}{\tau_1}}}{m/\tau_1} + \beta_2 \left(\frac{1 - e^{-\frac{m}{\tau_1}}}{m/\tau_1} - e^{-\frac{m}{\tau_1}}\right) + \beta_3 \left(\frac{1 - e^{-\frac{m}{\tau_2}}}{m/\tau_2} - e^{-\frac{m}{\tau_2}}\right) \qquad (3.6)
\end{aligned}$$

Now we have a model for the zero coupon interest rate as a function of maturity; the term structure. To find an equation for the term structure in the Nelson-Siegel model, just set $\beta_3 = 0$ in (3.6)

$$r(m) = \beta_0 + \beta_1 \frac{1 - e^{-\frac{m}{\tau_1}}}{m/\tau_1} + \beta_2 \left(\frac{1 - e^{-\frac{m}{\tau_1}}}{m/\tau_1} - e^{-\frac{m}{\tau_1}} \right) \tag{3.7}$$

We would like to be able to interpret the functional form in (3.7). In order to do so, note first that $\lim_{m\to\infty} r(m) = \beta_0$. The effect of β_0 on the yield curve is therefore permanent, and hence, we can interpret β_0 as the long-run component of the yield. All other things being equal, the long-term yield will approach β_0 as the maturity approaches infinity. This is also seen from the illustration in Figure 3.1, where it is seen that all of the illustrated shapes of Nelson-Siegel yield curves converge towards their β_0. The speed of convergence is primarily determined by the parameter τ_1.

Next, we want to investigate the short-term effect. We therefore take the limit of (3.7) as the maturity approaches zero. By use of l'Hôpital's rule[34], we obtain

$$\begin{aligned}\lim_{m\to 0} r(m) &= \beta_0 + \beta_1 \cdot \frac{0 - (-\frac{1}{\tau_1}) \cdot e^{-\frac{0}{\tau_1}}}{\frac{1}{\tau_1}} + \beta_2 \cdot \left(\frac{0 - (-\frac{1}{\tau_1}) \cdot e^{-\frac{0}{\tau_1}}}{\frac{1}{\tau_1}} - \lim_{m\to 0}(e^{-\frac{m}{\tau_1}}) \right) \\ &= \beta_0 + \beta_1 \end{aligned} \tag{3.8}$$

Hence, we can say that β_1 is a short-run component, since it starts out having full impact, but it declines to zero with increasing maturity. What remains is the β_2-part of (3.7). Note that the term involving β_2 in (3.7) starts out at zero, and also decreases to zero as m gets large. Therefore, it is fair to say that β_2 depicts a medium-run component of the term structure. Hence, we have an equation for the term structure with three different terms, and we can interpret the various components of the term structure as short-, medium- and long-run components.

The next step is of course to find a way to estimate the parameters, $\beta_0, \beta_1, \beta_2, (\beta_3), \tau_1$ (and τ_2) in this model. There are multiple ways to do this; non-linear least squares, maximum likelihood and generalized method of moments are the most obvious suggestions. Both Svensson (1994) and BIS (1999), however, note that a decision more important than the choice of optimization method, is the decision of whether to minimize the (sum of squared) price errors or yield errors. BIS (1999) argues that it makes most sense to minimize yield errors, if the aim of the estimation exercise is the term structure itself, and

[34]L'Hôpital's rule states that if $f(a) = g(a) = 0$ and $g'(a) \neq 0$, then $lim_{x\to a} \frac{f(x)}{g(x)} = \frac{f'(a)}{g'(a)}$. See e.g. Sydsæter (2000), p. 221.

not bond prices. On the other hand, they recognize that it is computationally easier to minimize price errors than yield errors. Their objection to minimizing price errors instead of yield errors, is that it leads to over-fitting of the long-term bond prices at the expense of the short-term bond prices.[35] Svensson (1994) notes that this is due to the insensitivity of short-term bond prices to interest rates. Therefore, it is advised to weight the observations with the inverse of their durations, when estimating the yield curve.[36] This is the approach that we use when estimating a term structure for the Danish mortgage bond market.

Usually the estimation is done by the least squares method, and this is exactly the path that we too will follow. The principle in the optimization algorithm is illustrated in Figure 3.2. In the first step, more or less arbitrary initial values of the parameters are assigned. These values are used to obtain a yield curve, which is subsequently used to calculate bond prices. Then the sum of the squared differences between observed and model prices is minimized by changing the parameters. These new parameter values are then used to obtain a new yield curve and the process continues until convergence is reached.

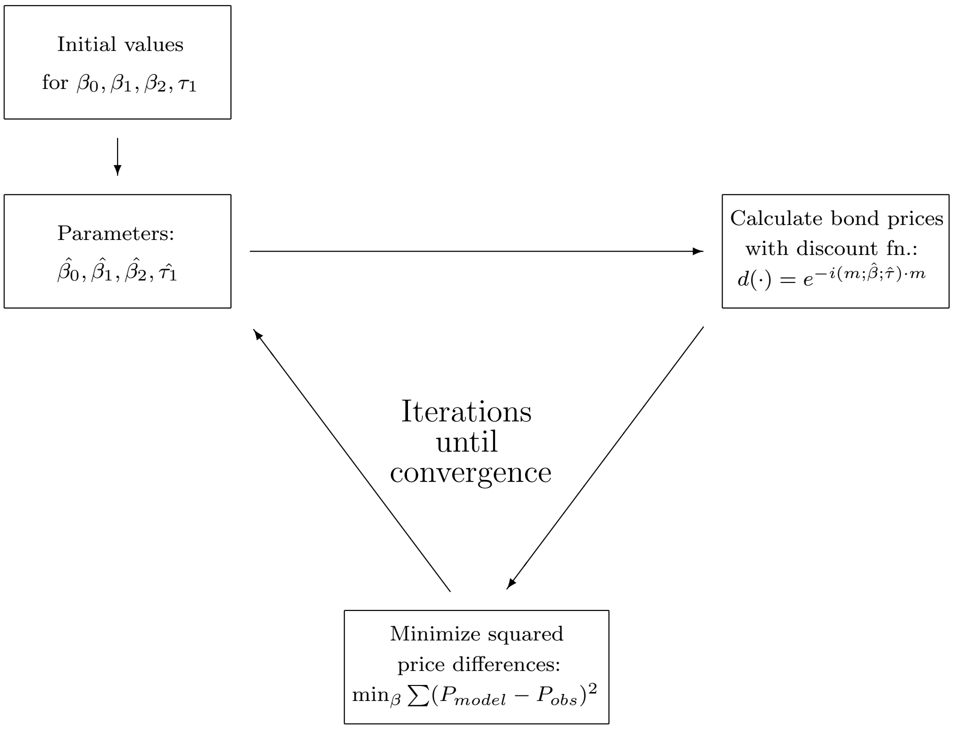

Figure 3.2: The estimation method in the Nelson-Siegel model

The main advantage of the least squares method is that it is relatively easy to implement, for instance in a somewhat sophisticated spreadsheet. We carry out the estimation

[35]BIS (1999), p. iii.
[36]The concept of duration is explained in detail in section 7.

in an Excel spreadsheet, where we have programmed the necessary functions in Visual Basic.[37]

3.1.2 Sample Selection

The next issue that we have to address, is which bonds to include in the estimation. Jakobsen (1992) notes that "*The ideal sample should consist of high liquidity bonds, distributed throughout the maturity spectrum and void any obstacles due to tax considerations or call features*". This sounds very simple and straightforward, but is very little so.

The first challenge is to find high liquidity bonds, distributed throughout the maturity spectrum. This is a hard task on the Danish market for mortgage bonds, as maturities are very unevenly distributed. However, this is a problem that could be overcome.

It is more problematic that the mortgage bonds included in the sample should void any tax- or prepayment option obstacles. We realize that ignoring tax considerations could be a source of problems in the estimation, but nevertheless we choose to disregard this issue. More importantly, it is hard to find mortgage bonds with long maturities that do not have an embedded prepayment option. Usually, this problem is dealt with such that one chooses the mortgage bonds with long maturities, on which the prepayment options are most out-of-the-money. Obviously this is done by choosing the callable bonds with the lowest coupon rate. This ensures that the value of the prepayment option is as small as possible, such that the value of the callable bond approaches the value of a non-callable bond with similar properties, cf. equation (1.1). If bonds with embedded prepayment options, of which the value is not negligible, are included in the sample, it would lead to estimation of interest rates that are too high. This is due to the fact that if the value of the prepayment option is larger than zero (as it is assumed), the value of the bond is deemed too low, leading to estimated yields that are too high. It is a fair point to say that this is is a fragile attempt to justify the use of callable bonds along with non-callable bonds in the estimation. However, it is the best readily available approximation we have, so we just have to bear in mind that this might be a cause of small biases in the final results.

The difficulties that arise when searching for a reliable sample for estimating a yield curve for Danish mortgage bonds, has led researchers to follow another path. It has become more and more common to use the swap-curve, i.e. a yield curve estimated on basis of quoted prices on interest rate swaps. Of course, it is best to use a yield curve

[37]The VBA functions are listed in the Appendix B.1, p. 135 and onwards.

that is estimated on basis of instruments that are fairly similar to the ones that one is trying to price, but on the other hand, if the obstacles to this approach are too large, it may be a good idea to use another yield curve that can be estimated with more ease and consistency over time. We choose to proceed with a yield curve estimated on basis of a sample of Danish mortgage bonds.

The next issue is whether the mortgage bonds in the sample should be from the same mortgage bank. From a theoretical point of view, this should clearly be the case, since differences in credit risk can lead to estimation biases. However, the credit risk on Danish mortgage bonds is regarded to be very small. This conclusion is supported by the fact that there has never, in the more than 200 years of mortgage financing in Denmark, been any defaults. All mortgage banks in Denmark have received a Standard & Poor's rating in the spectrum AA-AAA.[38] Since the credit risk is regarded to be very small, the *difference* in credit risk between different mortgage banks should also be small, and indeed negligible. The conclusion must be that if a satisfactory number of mortgage bonds from *one* mortgage bank exists, satisfying all other demands, one should definitely use these. If this is not the case, it does not constitute a major problem to include mortgage bonds from other mortgage banks. In our case, we do not encounter difficulties selecting a sample of mortgage bonds from the same mortgage bank.

Last, but not least, the sample should preferably not contain any foreign exchange rate risk. Even though the Danish central bank follows a policy of fixed exchange rates towards the Euro, bonds issued in Euro should not be included in the estimation, since there is a certain exchange rate risk on these bonds, which is not easily accounted for separately. Taking all these factors into consideration leads us to choose a sample of mortgage bonds as indicated in Table 3.1.[39]

Conducting the least squares estimation with the bonds in Table 3.1, and weighting each bond with the inverse of its duration, yields the parameter estimates in the Nelson-Siegel model as shown in Table 3.2.

The Nelson-Siegel yield curve estimated here is of a particularly simple form, since the medium-term component is very close to zero. The nicely shaped monotonous Nelson-

[38] Realkreditrådet (2005), p. 25. For more on Standard and Poor's rating methods and classifications, see www.standardandpoors.com.

[39] Please note that the durations in this table are calculated as $\sum_{t=1}^{T} \frac{PV(CF_t)\cdot t}{P}$ (see e.g. Grinblatt & Titman (2002) p. 826). This version of the duration is hardly suitable for any analysis, and even less for investment strategies. For the purpose at hand, namely weighting observations, we can use them without too much concern.

Issuer	Coupon	Maturity	Terms per year	Isin	Callable	Outst. Amount	Price	Duration
RD	4	2006	1	DK0009261323	No	124,923	100.160	0.10
RD	2	2006	1	DK0009270746	No	10,193	99.961	0.10
RD	4	2007	1	DK0009261240	No	32,338	101.286	1.07
RD	2	2007	1	DK0009270829	No	15,480	99.102	1.08
RD	4	2008	1	DK0009261166	No	28,891	101.981	1.99
RD	2	2008	1	DK0009270902	No	4,348	97.967	2.04
RD	4	2009	1	DK0009262131	No	8,064	102.591	2.88
RD	2	2009	1	DK0009271041	No	5,047	96.606	2.99
RD	4	2010	1	DK0009262990	No	6,690	102.825	3.74
RD	2	2010	1	DK0009271124	No	2,590	95.217	3.91
RD	4	2015	1	DK0009272015	No	623	102.664	7.57
RD	3	2028	4	DK0009274227	Yes	1,093	91.006	16.19
RD	4	2035	4	DK0009270233	Yes	42,233	95.182	17.12
RD	4	2038	4	DK0009274300	Yes	5,323	94.370	17.87

Note: Prices and liquidities as of November 21, 2005. Outstanding are measure in DKK mn.
Source: Copenhagen Stock Exchange

Table 3.1: Mortgage bond sample used in yield curve estimation

Parameter	**Estimate**
β_0	4.44 %
β_1	-2.04 %
β_2	0.00 %
τ_1	3.49

Table 3.2: Nelson-Siegel parameter estimates of the current yield curve (as of Nov 21, 2005)

Siegel fitted yield curve based on the parameter estimates shown in Table 3.2 is illustrated in Figure 3.3.

We will leave the initial yield curve shown in Figure 3.3 here for a moment, but make use for it in section 3.4, where we apply the term structure model, which we derive in the coming section.

3.2 Modelling the Term Structure

We now turn the view towards how to model the evolution of the term structure. A term structure model is a model, which – given a starting point – dictates the evolution of the yield curve. When choosing a model for the term structure of interest rates, we wish to use a model that fits data sufficiently, but also a model, which, for pedagogical reasons, is analytically tractable. We furthermore limit ourselves to looking at one-factor models, as

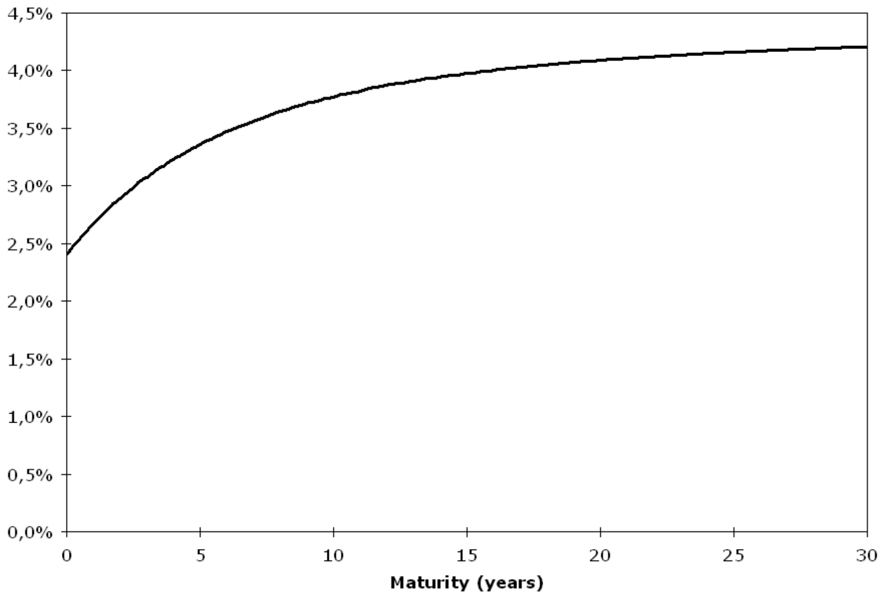

Figure 3.3: Nelson-Siegel fitted yield curve as of Nov 21, 2005

we want to obtain a fairly parsimonious model. By using a one-factor model, we assume that a single state variable summarizes all relevant information for pricing our interest rate dependent asset. As opposed to one-factor models, multi-factor models allow for more than one stochastic component to help explain the evolution in the term structure. However, the computational costs of using multi-factor models is extensive, and little is gained in the effort to understand the term structure of interest rates. Multi-factor models are thus out of the scope of this thesis, but such models can possibly provide the modeler with a better description of the term structure.

We initially provide a literature review to motivate our choice of model. Historically, the development in interest rates has been modelled as a stochastic differential equation (SDE), and the vast majority of these models are special cases of Itô processes. We discuss the different features embedded in these models, and subsequently choose a model for further use in our pricing model.

Early contributions focused on modelling the interest rate in a general equilibrium setting. Examples of such models are proposed in Merton (1973), Vasicek (1977) and Cox, Ingersoll & Ross (1985). Equilibrium models seek to explain how general underlying economic variables influence the interest rate. Hence, one obtains a present term structure as an output of the model, based on assumptions about risk preferences and supply and demand relationships between bonds and other assets etc.

More recent contributions have instead modelled the term structure in an arbitrage-free set-up. Examples of such models are proposed in Ho & Lee (1986), Hull & White (1990*a*) and Heath, Jarrow & Morton (1992). As opposed to equilibrium models, these models use today's term structure as an input. The modeler estimates today's term structure using a statistical model, aiming primarily at a satisfying fit to observed asset prices, while paying less attention to explanatory power. This is exactly what we did in section 3.1 using the Nelson-Siegel model. Governed by structural assumptions, the modeler then uses the present yield curve to determine the future average path taken by the instantaneous interest rate. We now present the evolution of the one-factor models as to shed light on the aspects that a modeler has to consider when choosing a model.

3.2.1 Examples of Models

Merton was among the first modern economists to formalize the term structure. He suggested an equilibrium model, in which the interest rate process under the Q-measure, could be described by an arithmetic Brownian motion[40] with both the drift term and the diffusion term being constant. Thus, the change in the interest rate can be written as

$$dr = \mu dt + \sigma dW^Q \tag{3.9}$$

If one solves for r, one relatively easily sees that the short interest rates are normal.[41] Models with this property are called Gaussian models. The Gaussian density function of r makes the model analytically tractable and provides us with a log-normal asset price. We return to this in section 3.2.2, where we solve such a model.

Gaussian models have the obvious flaw of assigning positive probabilities to negative interest rates. It is an undesirable property as the term structure is most often modelled in nominal terms, and negative nominal rates would imply possible arbitrage.[42] Furthermore, negative nominal interest rates are rarely observed.

Another significant shortfall of the Merton model is that the proposed SDE does not prevent the interest rate from drifting off to either positive or negative infinity.[43] Also, there exists no intuitive argument why the interest rate should have a non-zero constant drift rate. History has shown that several economic variables including interest rates are mean-reverting – that is they have a steady state level that they tend to be drawn towards.

[40]See Rendleman & Bartter (1980) for a model using a *geometric* Brownian motion.

[41]Duffie (2001), p. 139.

[42]We refer to Duffie (2001) p. 140 for a discussion of this special kind of arbitrage.

[43]In fact, the zero coupon bond price implied by the Merton model converges to positive infinity as the maturity goes to infinity. See Cairns (2004) p. 76.

Hence, mean-reversion is generally thought to be a desirable property of a term structure model. This led Vasicek to develop a model including mean-reversion. The Vasicek SDE is also called an Ornstein-Uhlenbeck process[44] and it can be written as

$$dr = \gamma(\theta - r)dt + \sigma dW^Q \tag{3.10}$$

where γ, θ and σ are strictly positive constants. It can be seen that the model implies mean reversion with γ being the parameter indicating the speed of reversion to the steady state equivalent martingale level θ. The Vasicek model is a valuable contribution to the field of research due to its mathematical convenience. However, as in the Merton model, the convenience comes at the expense of non-negative nominal interest rates. Combined with a generally poor fit to empirical evidence, the model is of limited use to practitioners, but it is often used for introductory academic purposes.

As it became increasingly clear that equilibrium models could not provide practitioners with a satisfying fit between model and observed interest rates, the modern class of term structure models, arbitrage-free models, was introduced by Ho & Lee (1986).

Ho & Lee introduced a simple model that extended the Merton model by including a time-dependent drift term.[45] Its SDE has the following representation

$$dr = \mu(t)dt + \sigma dW^Q \tag{3.11}$$

Except for a superior fit (by using $\mu(\cdot)$ to fit the current term structure), it does not provide a solution for the flaws of Vasicek. It allows for negative interest rates, and in addition to this, it omits mean reversion. The Ho-Lee model has no broad application today as more advanced arbitrage-free models have proven to be superior.

The first tractable model to ensure non-negative interest rates was the Cox, Ingersoll & Ross (1985) (CIR) model, which also was introduced to deal with the shortcomings of the Vasicek model. It extended the Vasicek model by including an interest rate dependent diffusion term, and its diffusion process can be written as

$$dr = \gamma(\theta - r)dt + \sigma\sqrt{r}dW^Q \tag{3.12}$$

where γ, θ and σ are positive real variables. It can thus be seen that $r = 0$ is a reflecting

[44]Dixit & Pindyck (1993) p. 74.

[45]Originally, Ho-Lee proposed a model using a binomial tree, which Dybvig (1997) and Jamshidian (1988) showed to have the SDE in (3.11) as a limit case.

barrier. If r reaches zero, the diffusion term also equals zero, and the drift term will take r into the strictly positive domain. Hence, it produces mean-reverting interest rates as the Vasicek model, and it, furthermore, restricts $r(t)$ to $\mathbb{R}_+$. However, the choice of $\sqrt{r}$ might seem somewhat arbitrary. Although empirical evidence (see e.g. Chan, Karolyi, Longstaff & Sanders (1992)) has found that interest rate volatility seems to be increasing in the interest rate, it does not suggest that a square root function is the exact dependence. The most important feature of the CIR model is, however, not the exact exponent of the interest rate, but that it restricts the interest rate to the positive domain. Despite having the desired properties for a term structure model, the CIR model cannot completely escape the curse of equilibrium models, which is an unsatisfying fit.

It should be clear from the model review so far that several models have been proposed – all of which slightly improves earlier contributions. However, not until 1992 did anyone propose a complete formalization of the term structure. In the seminal paper Heath, Jarrow & Morton (1992), the authors present a general multi-factor framework for forward rates. This model has the previously mentioned models as special cases. Unfortunately, but not surprisingly, one had to leave the Gaussian model class and thus analytical tractability, and use of the general HJM model therefore implies that valuation is carried out using numerical methods. We have therefore chosen to apply the Hull-White model, which is a special case of the HJM setting. More importantly, the Hull-White model is used to a great extent by practitioners, which can be seen as an indication of its validity.

3.2.2 Hull-White Model

The next step in the venture of setting up a term structure model is to find expressions for bond and derivatives prices in the Hull-White framework, as this enables us to calibrate the model parameters later in section 3.3.

Hull & White (1990*a*) proposed an arbitrage-free model with the following SDE

$$dr = [\theta(t) - a(t)r]dt + \sigma(t)dW^Q$$

Since this SDE has a deterministic volatility, we know that it is also a Gaussian model. It can be seen that the Hull-White SDE allows for a high degree of freedom, since it allows for both a time-dependent drift term and a time-dependent diffusion term. However, we choose to keep the mean-reversion parameter and the volatility constant. Our choice is mainly motivated by the fact that it is noted in Hull & White (1994) that by allowing a and σ to be time-dependent, potentially little is gained. If one wants to put emphasis on the volatility structure, a multi-factor should be applied rather than settling for an initially estimated volatility structure. Hull & White note that the additional technical complexity of allowing for a time-dependent diffusion term, is only a slight improvement compared to a constant diffusion term. The inclusion of a time-dependent diffusion term can cause a significant bias in the long end of the curve, if the volatility structure changes considerably. For reasons of simplicity, we choose to keep the parameters constant. Hence, we use the following SDE[46]

$$dr = [\theta(t) - ar]dt + \sigma dW^Q \tag{3.13}$$

where a and σ are constants and $\theta(\cdot)$ is a deterministic function of time. Under the Q-measure, the short interest rate reverts to $\frac{\theta(t)}{a}$ with $\frac{1}{a}$ being the reversion speed. It is easily seen that this simplified Hull-White model is equivalent to the Vasicek model with a time-dependent mean reversion level or the Ho-Lee model with mean reversion. The Hull-White model also assigns positive probability to negative interest rates. This probability can relatively easily be calculated, and it is of course increasing in σ and decreasing in μ.[47]

We now solve the model, such that we obtain the specific pricing formulas for the Hull-White model based on the results from the pricing section. We know from section 2 that under the Q-measure, the Hull-White representation must satisfy the general term structure equation from Result 2.1. Replacing the general functions with those of the Hull-White model provides us with

$$F_t + [\theta(t) - ar(t)]F_r + \frac{1}{2}\sigma^2 F_{rr} - r(t)F = 0 \tag{3.14}$$

$$F(T;T) = 1 \tag{3.15}$$

A solution to this differential equation is known as an affine price equation.[48] We thus continue with F having the general form of an affine price function, and we write it as

$$F(t;T) = e^{A(t,T) - B(t,T)r} \tag{3.16}$$

We now move on to solve for the price coefficient functions $A(\cdot)$ and $B(\cdot)$ such that we obtain the exact form for $F(\cdot)$. We insert the relevant derivatives of equation (3.16) into

[46] Though a simplified version, we henceforth refer to it as the Hull-White model throughout this thesis.
[47] See Brigo & Mercurio (2001), p. 65.
[48] Dana & Jeanblanc (2003), p. 172.

(3.14) to get

$$\begin{aligned}\underbrace{(A_t - B_t r)F}_{F_t} + (\theta(t) - ar)\underbrace{(-BF)}_{F_r} + \frac{1}{2}\sigma^2 \underbrace{B^2 F}_{F_{rr}} - rF &= 0 \quad \Rightarrow \\ \left[(A_t - \theta(t)B + \frac{1}{2}\sigma^2 B^2) - (B_t - aB - 1)r\right] F &= 0 \qquad (3.17)\end{aligned}$$

This provides us with the conditions under which (3.16) is a solution. Equation (3.17) must be satisfied for all maturities and since the interest rate is independent of T, the coefficient of r must be zero; that is

$$B_t - aB - 1 = 0 \qquad (3.18)$$

which in turn gives us

$$A_t - \theta(t)B + \frac{1}{2}\sigma^2 B^2 = 0 \qquad (3.19)$$

We know from the boundary condition that at time T, the asset price must equal 1 independently of the realized interest rate at time T. By setting the coefficient of r equal to zero we obtain

$$A(T,T) = B(T,T) = 0 \qquad (3.20)$$

We can now derive the expressions for $A(\cdot)$ and $B(\cdot)$ using (3.18)-(3.20).

Equation (3.18) is easily recognizable as a linear ordinary differential equation in t (for a fixed maturity), which we solve as follows

$$\begin{aligned}B(t,T) &= Ce^{at} + \frac{1}{a} \Rightarrow \\ B(t,T) &= \frac{1}{a}\left(1 - e^{-a(T-t)}\right) \qquad (3.21)\end{aligned}$$

We have solved for the constant C using the boundary condition. Having calculated $B(\cdot)$ we can now solve for $A(\cdot)$. Equation (3.19) can be rearranged and subsequently integrated into

$$\begin{aligned}A_t &= \frac{1}{2}\sigma^2 B^2 - \theta(t)B \Rightarrow \\ A(t,T) &= \int_t^T \left(\frac{1}{2}\sigma^2 B^2(s,T) - \theta(s)B(s,T)\right) ds \qquad (3.22)\end{aligned}$$

We have thus derived the price coefficients, $A(\cdot)$ and $B(\cdot)$, as functions of a, σ and $\theta(t)$.

This is where the Hull-White model extends the Vasicek model. In the Vasicek model, θ would be a constant and we would then have completed the calculation of our price expression. However, in the Hull-White model θ is a time-dependent function, which we now determine such that the model fits the initial term structure.

We choose $\theta(\cdot)$ such that the theoretical prices $\{p(0,T); T > 0\}$ fit the observed prices $\{\hat{p}(0,T); T > 0\}$ and therefore the initial term structure. It is more convenient to fit prices by using the forward rate, which contracted at time t with maturity T, is defined as[49]

$$f(t,T) = -\frac{\partial \ln p(t,T)}{\partial T} \tag{3.23}$$

From (3.16) it readily follows that

$$f(0,T) = B_T(0,T)r(0) - A_T(0,T) \tag{3.24}$$

where

$$\begin{aligned} B_T(0,T) &= \frac{\partial}{\partial T}\left(\frac{1}{a}(1-e^{-a(T-t)})\right) \\ &= e^{-aT} \end{aligned} \tag{3.25}$$

$$\begin{aligned} A_T(0,T) &= \frac{\partial}{\partial T}\int_0^T \frac{1}{2}\sigma^2 \underbrace{\frac{1}{a^2}(1-e^{-a(T-s)})^2}_{B^2(s,T)} ds - \frac{\partial}{\partial T}\int_0^T \theta(s) \underbrace{\frac{1}{a}(1-e^{-a(T-s)})}_{B(s,T)} ds \\ &= \frac{\sigma^2}{2a^2}(1-e^{-aT})^2 - \int_0^T e^{-a(T-s)}\theta(s)ds \end{aligned} \tag{3.26}$$

By inserting (3.25) and (3.26) into (3.24), we obtain

$$f(0,T) = e^{-aT}r(0) + \int_0^T e^{-a(T-t)}\theta(s)ds - \frac{\sigma^2}{2a^2}(1-e^{-aT})^2 \tag{3.27}$$

We then solve (3.27) for any $T \geq 0$ given the observed initial forward term structure, that is by solving

$$\hat{f}(0,T) = e^{-aT}r(0) + \int_0^T e^{-a(T-t)}\theta(s)ds - \frac{\sigma^2}{2a^2}(1-e^{-aT})^2 \tag{3.28}$$

[49] Björk (1998) p. 222.

for $\theta(\cdot)$. To solve this equation we apply a trick by writing

$$\hat{f}(0,T) = x(T) - g(T) \tag{3.29}$$

where x and g are defined as follows

$$\dot{x} = -ax(t) + \theta(t), \quad x_0 = r_0 \tag{3.30}$$

$$g(t) = \frac{\sigma^2}{2} B^2(0,t) \tag{3.31}$$

Now rearranging (3.30), inserting that $\hat{f}_T(0,T) = \dot{x}(T) - \dot{g}(T)$, and subsequently inserting (3.29) gives us

$$\begin{aligned} \theta(T) &= \dot{x}(T) + ax(T) \\ &= \hat{f}_T(0,T) + \dot{g}(T) + ax(T) \\ &= \hat{f}_T(0,T) + \dot{g}(T) + a\left(\hat{f}(0,T) + g(T)\right) \end{aligned} \tag{3.32}$$

Thus, if we choose $\theta(\cdot)$ according to (3.32), we obtain a term structure that implies a perfect fit between our model-predicted current prices ($p(0,T)$) and observed current prices ($\hat{p}(0,T)$) for any $T \geq 0$. We now insert (3.32) into (3.22)

$$\begin{aligned} A(t,T) &= \int_t^T \left(\frac{1}{2}\sigma^2 B^2(s,T) - \left[\hat{f}_T(0,s) + \dot{g}(s) + a\left(\hat{f}(0,s) + g(s)\right)\right] B(s,T)\right) ds \\ &= B(t,T)\hat{f}(0,t) - \frac{\sigma^2}{4a} B^2(t,T)(1 - e^{-2at}) + \ln\left(\frac{p(0,T)}{p(0,t)}\right) \end{aligned} \tag{3.33}$$

Hence, by substituting (3.33) into (3.16), we obtain the theoretical bond price as a function of $B(t,T)$ as follows

$$\begin{aligned} p(t,T) &= F(t,T) \\ &= e^{A(t,T) - B(t,T)r(t)} \\ &= \frac{p(0,T)}{p(0,t)} \exp\left(B(t,T)\hat{f}(0,t) - \frac{\sigma^2}{4a} B^2(t,T)(1 - e^{-2at}) - B(t,T)r(t)\right) \end{aligned}$$

We have thus obtained the bond price. For reasons of completeness, we state the price of a zero coupon bond with a principal of \$$L$.

Result 3.1 ***Hull-White Zero Coupon Bond Price***

When using the Hull-White term structure model, the price of a zero coupon bond paying $L at time T is

$$p(t,T,L) = L\cdot\frac{p(0,T)}{p(0,t)}\exp\left(B(t,T)\hat{f}(0,t) - \frac{\sigma^2}{4a}B^2(t,T)(1-e^{-2at}) - B(t,T)r(t)\right) \quad (3.34)$$

where $B(t,T) = \frac{1}{a}\left(1-e^{-a(T-t)}\right)$.

This initially completes the model, as we now have an expression for the bond price in the Hull-White model. However, when calibrating the model in section 3.3, we use prices of derivatives and we therefore need the theoretical terms of such assets. Hence, we now solve for option prices in the Hull-White model.

Recall that according to (2.33), the general call option price on a zero coupon bond could be stated as

$$ZBC(t,T_1,T_2,K,L) = L\cdot p(t,T_2)Q^{T_2}\{L\cdot p(T_1,T_2) > K\} - K\cdot p(t,T_1)Q^{T_1}\{L\cdot p(T_1,T_2) > K\}$$

To facilitate computability of the probabilities, we require that the numeraire process (M) has a deterministic volatility. Recall the definition of M is

$$\begin{aligned} M(t) &= \frac{p(t,T_2)}{p(t,T_1)} \\ &= e^{A(t,T_2)-A(t,T_1)-[B(t,T_2)-B(t,T_1)]r(t)} \end{aligned}$$

where we have inserted (3.16). As the volatility term is unaffected by a change of measure, it suffices to check whether the volatility is deterministic under one measure. We verify that it is deterministic under the Q-measure. Applying Itô's lemma gives us the Q-dynamics of $M(\cdot)$ as

$$\begin{aligned} dM(t) &= \left(\underbrace{\frac{\partial M}{\partial t} + (\theta(t)-ar)\frac{\partial M}{\partial r} + \frac{1}{2}\sigma\frac{\partial^2 M}{\partial r^2}}_{m(t)}\right)dt + \sigma\frac{\partial M}{\partial r}dW \\ &= M(t)\left(m(t)dt - \sigma[B(t,T_2)-B(t,T_1)]dW\right) \\ &= M(t)\left(m(t)dt + \underbrace{\frac{\sigma}{a}e^{at}(e^{-aT_2}-e^{-aT_1})}_{\sigma_M}dW\right) \end{aligned}$$

The volatility term, σ_M, is indeed deterministic. We can thus apply the general option pricing formula stated above. However, we need to calculate the exact expression for Σ^2 as this is model specific.

$$\begin{aligned}
\Sigma^2 &= \int_t^{T_1} \frac{\sigma^2}{a^2} e^{2as}(e^{-aT_2} - e^{-aT_1})^2 ds \\
&= \frac{\sigma^2}{2a^3}(e^{-2aT_2} + e^{-2aT_1} - 2e^{-aT_2-aT_1})(e^{2aT_1} - e^{2at}) \\
&= \frac{\sigma^2}{2a^3}(1 - e^{-2a(T_1-t)})(1 - e^{-a(T_2-T_1)})^2 \qquad (3.35)
\end{aligned}$$

Finally, we write the price of a European call option on a zero coupon bond using the Hull-White term structure model as

$$ZBC(t, T_1, T_2, K, L) = L \cdot p(t, T_2)\Phi(h) - K \cdot p(t, T_1)\Phi(h - \Sigma) \qquad (3.36)$$

where

$$h = \frac{1}{\Sigma} \ln \left[\frac{L \cdot p(t, T_2)}{K \cdot p(t, T_1)} \right] + \frac{\Sigma}{2} \qquad (3.37)$$

$$\Sigma = \frac{\sigma}{a}(1 - e^{-a(T_2-T_1)}) \sqrt{\frac{(1 - e^{-2a(T_1-t)})}{2a}} \qquad (3.38)$$

For reasons of convenience and consistency, we also state the price of a European put option on a zero coupon bond at this point. We look at a put option with similar characteristics as the call option above. The price of such a put option on a zero coupon bond is given by

$$ZBP(t, T_1, T_2, K, L) = K \cdot p(t, T_1)\Phi(-h + \Sigma) - L \cdot p(t, T_2)\Phi(-h) \qquad (3.39)$$

where h and Σ are defined above.

Now that we have obtained pricing formulas under the Hull-White model, we move on to calibrate the model in the next section. In particular, we show how to estimate the mean-reversion and volatility parameters using the option pricing formulas just developed.

3.3 Volatility and Model Calibration

The next issue is to estimate the parameters in the Hull-White model. Before we do this, however, we note that there is a minor problem using a Nelson-Siegel yield curve along with the Hull-White model. The problem is, as Björk & Christensen (1999) showed, that the Hull-White model is inconsistent with the Nelson-Siegel yield curve family. This is to be understood in the way that the Hull-White model can of course accommodate a Nelson-Siegel yield curve as an input, but the Hull-White model will in general produce forward interest rate term structures that are not representable by the Nelson-Siegel functional form.[50] The Nelson-Siegel family is, in a sense, too small to capture all kinds of yield curves that can be produced by the Hull-White model in future periods. Björk & Christensen (1999) proceed to show that it only requires a slight modification of the Nelson-Siegel functional form in order to enable the functional form to accommodate all possible outcomes of the Hull-White model. We do not put too much emphasis on this objection to the combination of the Hull-White model and the Nelson-Siegel family of yield curves, since the only use for the Nelson-Siegel yield curve for us, is to be able to make a better estimate of the current term structure than simple bootstrapping. Though theoretically interesting, the fact that the future generated forward yield curves, produced by the Hull-White model, are not representable by a Nelson-Siegel functional form, is not such a great concern to us in the present context.

For the purpose of calibrating the Hull-White model, we need a list of so-called calibrating instruments, which are securities that can be valued inside the Hull-White framework. These securities are normally chosen to be so-called *caps* or *floors*, but swaptions could also be used for this purpose.[51] Hence, in the following we develop pricing formulas for such instruments, such that we can calibrate the model by matching observed and model prices. Before we do this, we will briefly go through the necessary concepts.

A *cap* is a financial instrument that is made to give insurance to a borrower against a rise in the interest rate, on a floating-rate loan.[52] Consider a loan of maturity T that is based on some floating interest rate, e.g. some CIBOR or LIBOR[53] rate. The interest rate on the loan is periodically reset to the underlying interest rate. The time between two resets is referred to as the *tenor*; we denote the reset dates by $t_0, t_1, \dots t_n$, where $t_n = T$ and $\delta_k = t_{k+1} - t_k$ is the (usually constant) tenor.

The idea of a cap is that if at a given reset date t_k, the interest rate underlying the floating-rate of the bond rises above a predetermined level, called the *cap rate*, the borrower still only has to pay an interest rate equal to the cap rate in the period between

[50] Björk & Christensen (1999), p. 338.

[51] Hull (2000), p. 593-594. Usually, caps and floors are used to calibrate models, except in case one uses a yield curve based on interest rate swaps, in which case it is more logical to use swaptions.

[52] This section is partly based on Hull (2000) and Björk (1998).

[53] Copenhagen InterBank Offered Rate and London InterBank Offered Rate, respectively.

t_k and t_{k+1}. Initially, we consider a product that only provides a cap on the interest rate on a loan in one period, between t_k and t_{k+1}. Such a product is called a *caplet*, and obviously a cap can be interpreted as a collection of caplets. The payoff to the holder of a caplet at time t_{k+1} with cap rate r_{cap} and underlying interest rate at r_k in the period between t_k and t_{k+1} is

$$\xi_{k+1}^{caplet} = L\delta_k \cdot \max\left[r_k - r_{cap}, 0\right] \tag{3.40}$$

where L is the principal and δ_k is the tenor. Equation (3.40) is, by definition, the value of a call option on the underlying interest rate at time t_k with payment at time t_{k+1}. Since, as we noted previously, a cap can be interpreted as a collection of caplets, we can also interpret a cap as a collection of call options on the underlying interest rate. In order to find the value of the payoff of a caplet at time t_k, we discount (3.40) with $\delta_k r_k$.

$$\begin{aligned}
\xi_k^{caplet} &= \frac{\xi_{k+1}^{caplet}}{1+\delta_k r_k} \\
&= \frac{L\delta_k}{1+\delta_k r_k} \cdot \max[r_k - r_{cap}, 0] \\
&= \max\left[\frac{L\delta_k r_k + L - L - L\delta_k r_{cap}}{1+\delta_k r_k}, 0\right] \\
&= \max\left[L - \frac{L(1+\delta_k r_{cap})}{1+\delta_k r_k}, 0\right] \qquad (3.41)
\end{aligned}$$

Note that $\frac{L(1+\delta_k r_{cap})}{1+\delta_k r_k}$ is the time t_k value of a zero coupon bond that pays off $L(1+\delta_k r_{cap})$ at time t_{k+1}. Therefore, we can see (3.41) as the value of a put option on a zero coupon bond with face value $L(1+\delta_k r_{cap})$ and strike price L. Hence, the cap can both be seen as a collection of call options on the underlying interest rate, and now also as a collection of put options on zero coupon bonds. This observation can be used to price caps in any pricing model that is able to price call and put options, and this is an observation that will come in handy shortly.

Similar to the concept of a cap, there is also another type of interest rate derivative known as a *floor*. The concept of a floor is that the seller can oblige himself to pay a certain minimum interest rate on a loan in case the underlying interest rate should fall below a certain level, the *floor rate*. Here, it is the holder of the bond that obtains insurance against adverse movements (from his point of view) in the underlying interest rate. The issuer of the bond (the borrower) obtains a premium for committing himself to pay a certain minimum interest rate. The pricing concepts are similar, and the floor can be seen as (i) a portfolio of put options on the underlying interest rate, or (ii) a portfolio of call options on zero coupon bonds. Each of these call options is known as a *floorlet*.

A floor and a cap can be combined to create a so-called *collar*, which ensures an interest rate between the floor rate and the cap rate. For instance, the collar can be created in such a way that the cap and the floor are balanced to make the combined derivative liquidity neutral, meaning that the premium of selling off the floor exactly equals the cost of buying the cap.[54] The concepts of a cap, a floor, and a collar are illustrated in Figure 3.4.

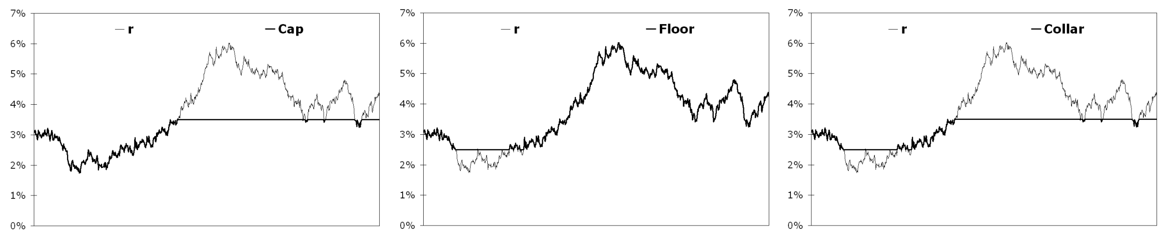

Figure 3.4: Interest payments with a cap, a floor, and a collar

Ultimately, a floating-rate loan can be transformed into a fixed-rate loan by buying a cap and selling a floor with the same strike (cap rate = floor rate). Obviously, this means that the price of such an arrangement has to equal the price of a swap that swaps floating-rate interest rate payments into fixed-rate interest payments with the same fixed interest rate as the cap/floor strike rate. Otherwise, arbitrage profit could be earned. This is the put-call parity of caps and floors:[55]

$$\text{cap price} - \text{floor price} = \text{swap price}$$

Now that we have found out that we can interpret a cap as a collection of put options on zero coupon bonds, we will use this observation by applying the pricing formula for put options in the Hull-White framework that we derived in section 3.2.2. Afterwards, we will do similarly for a floor, using the pricing formula for a call option on a zero coupon bond.

The time t-value of a put option with strike price K and maturity T_1 on a zero coupon bond maturing at time T_2 with principal L was shown to be

$$ZBP(t, T_1, T_2, K, L) = K \cdot p(t, T_1) \cdot \Phi(-h + \Sigma) - L \cdot p(t, T_2) \cdot \Phi(-h) \tag{3.42}$$

[54] We return to a brief discussion of the use of these instruments (caps, floors, and collars) in the Danish mortgage bond market in section 9.

[55] For a given *equal* interest rate, obviously.

while the time-t value of a call option with strike price K and maturity T_1 on a zero coupon bond maturing at time T_2 with principal L was shown to be

$$ZBC(t, T_1, T_2, K, L) = L \cdot p(t, T_2) \cdot \Phi(h) - K \cdot p(t, T_1) \cdot \Phi(h - \Sigma) \tag{3.43}$$

where $p(\cdot)$ is the price of the bond according to the Hull-White model, $\Phi(\cdot)$ is the cumulative standardized normal distribution, and h and Σ are given by

$$\begin{aligned} h &= \frac{1}{\Sigma} \cdot \ln\left[\frac{L \cdot p(t, T_2)}{K \cdot p(t, T_1)}\right] + \frac{\Sigma}{2} \\ \Sigma &= \frac{\sigma}{a}\left(1 - e^{-a(T_2 - T_1)}\right) \cdot \sqrt{\frac{1 - e^{-2a(T_1 - t)}}{2a}} \end{aligned}$$

As mentioned, the value of a cap at time t with reset dates $\{t_k\}_{k=1}^n$, principal of the bond L, and cap rate r_{cap}, leading to a strike price of $\frac{L}{1+r_{cap}\delta_k}$, can be calculated as a sum of the values of a collection of put options on zero coupon bonds, and the value of a cap is therefore given by

$$\text{Cap}\left(t, \{t_k\}_{k=1}^n, L, \frac{L}{1 + r_{cap}\delta_k}\right) = \sum_{k=1}^{n} \left((1 + r_{cap}\delta_k) \cdot ZBP\left(t, t_{k-1}, t_k, \frac{L}{1 + r_{cap}\delta_k}, L\right)\right) \tag{3.44}$$

If we insert (3.42), we get the following

$$\begin{aligned} \text{Cap}(\cdot) &= \sum_{k=1}^{n} \left[(1 + r_{cap}\delta_k) \cdot \left(\frac{L}{1 + r_{cap}\delta_k} \cdot p(t, t_{k-1}) \cdot \Phi(-h_k - \Sigma_k) - L \cdot p(t, t_k) \cdot \Phi(-h_k)\right)\right] \\ &= L \cdot \sum_{k=1}^{n} \left[p(t, t_{k-1}) \cdot \Phi(-h_k - \Sigma_k) - (1 + r_{cap}\delta_k) \cdot p(t, t_k) \cdot \Phi(-h_k)\right] \end{aligned} \tag{3.45}$$

where

$$\begin{aligned} h_k &= \frac{1}{\Sigma_k} \cdot \ln\left[\frac{L \cdot p(t, t_k)}{p(t, t_{k-1}) \cdot \frac{L}{1+r_{cap}\delta_k}}\right] + \frac{\Sigma_k}{2} \\ &= \frac{1}{\Sigma_k} \cdot \ln\left[\frac{p(t, t_k) \cdot (1 + r_{cap}\delta_k)}{p(t, t_{k-1})}\right] + \frac{\Sigma_k}{2} \qquad (3.46) \\ \Sigma_k &= \frac{\sigma}{a}\left(1 - e^{-a(t_k - t_{k-1})}\right) \cdot \sqrt{\frac{1 - e^{-2a(t_{k-1} - t)}}{2a}} \qquad (3.47) \end{aligned}$$

The derivation principle of the pricing formula for a floor is obviously very similar to deriving the pricing formula of a cap. The value of a floor at time t with reset dates

$\{t_k\}_{k=1}^n$, principal of the bond L and cap rate r_{cap}, leading to a strike price of $\frac{L}{1+r_{cap}\delta_k}$, is hence given as the sum of the values of a collection of call options on zero coupon bonds:

$$\text{Floor}\left(t, \{t_k\}_{k=1}^n, L, \frac{L}{1+r_{cap}\delta_k}\right) = \sum_{k=1}^{n}\left((1+r_{cap}\delta_k)\cdot ZBC(t, t_{k-1}, t_k, \frac{L}{1+r_{cap}\delta_k}, L)\right)$$

Plugging in $ZBC(\cdot)$ and rearranging terms yields

$$\text{Floor}(\cdot) = L\cdot\sum_{k=1}^{n}\left[(1+r_{cap}\delta_k)\cdot p(t,t_k)\cdot\Phi(h_k) - p(t,t_{k-1})\cdot\Phi(h_k-\Sigma_k)\right] \tag{3.48}$$

where h_k and Σ_k are given by (3.46) and (3.47). Since we now have analytical formulas for pricing caps and floors, we can proceed to calibrate the model by fitting the model prices of caps and floors as given by the expressions in (3.45) and (3.48) to observed market prices. To be able to do this consistently, we have to set up a goodness of fit measure. An immediate choice is to minimize the sum of the squared errors between observed prices p_j and the model calculated prices $\hat{p_j}$ for the $j = 1 \ldots m$ cap and/or floor prices:

$$\min_{a,\sigma}\sum_{j=1}^{m}(p_j-\hat{p_j})^2 \tag{3.49}$$

This is particularly straightforward in our case, where neither a nor σ is a function of time. Had this not been the case, it would have been necessary to divide the maturity span into smaller segments for the parameter that is allowed to change,[56] or to specify deterministic functional form(s) for $a(t)$ and/or $\sigma(t)$. In order to ensure that the function(s) that is/are time dependent do(es) not change dramatically over time, so-called penalty functions are often employed. However, in our case, we can proceed directly to make the calibration of the model, since we have assumed that neither of the parameters are time-dependent.

To carry out the calibration of the model, we need market prices for caps and/or floors. This, however, poses a new challenge. Prices of caps and floors are normally not quoted in terms of direct prices; instead they are quoted by the use of implied volatilities. Implied volatilities means the implied volatilities of the underlying interest rate. These volatilities are, however, model dependent. The market standard is to quote the implied volatilities under the assumption of a log-normally distributed interest rate. This is precisely the

[56]The procedure works the same way if both parameters are allowed to be time dependent, but of course the estimation will be conducted with more uncertainty (higher standard errors) if both a and σ are allowed to change.

assumption underlying the Black-76 model[57], which is why the implied volatilities are usually denoted Black-76 volatilities. One uses the value of a caplet in the Black-76 model, and one inserts the implied volatility σ_k as quoted in the market. This gives a market price, which we can fit to the model price from the Hull-White model.

Again, this raises another question; is the cap volatility assumed to be constant for all the embedded caplets? The answer is usually yes, but not always. If the volatility is assumed to be constant for all the embedded caplets of a cap, we denote the implied volatility, the *flat volatility*. If not, the cap volatility is denoted the *spot volatility*. The relationship between flat volatilities and spot volatilities is actually analogue to the relationship between yield to maturity and zero coupon spot interest rates. So, the flat volatility is a weighted average volatility, with some of the same shortcomings as is the case when using yield to maturity as the interest rate. Nevertheless, it is market standard to quote the prices on the cap market as implied *flat* volatilities. This is important to note, when conducting the estimation.

The next problem that arises, is that cap volatilities for DKK are only quoted for maturities up to ten years.[58] If we want to include volatilities in the calibration that have a longer maturity than ten years, we have to decide which cap volatilities we want to use, as directly observable DKK volatilities are not available. It is industry practice to use the Euro volatilities as guideline, since Euro volatilities are also available for maturities of 15 and 20 years, and due to the fixed rate regime in Denmark towards the Euro. However, we can of course not just use the Euro cap volatilities for maturities of 15 and 20 years along with DKK cap volatilities for maturities of up to 10 years, without considering what kind of correction of the Euro cap volatilities should be applied. In general, the DKK cap volatilities are higher than the Euro cap volatilities, primarily due to the existence of a liquidity premium and foreign exchange rate risk. Higher uncertainty on the underlying factors of the interest payments, obviously makes the expected volatility of interest rates higher on DKK, which is why the implied volatilities for DKK are higher than for Euro. Again, it is industry practice to use the Euro cap volatilities added one percentage point for maturities not directly available in DKK. Another approach could be to calculate the average markup for DKK cap volatilities compared to Euro cap volatilities for maturities up to 10 years, and scale the Euro cap volatilities for maturities of 15 and 20 years up with this factor. Any of these methods are of course only reliable if there is a somewhat stable relationship between DKK and Euro cap volatilities, either in absolute or in relative

[57] See e.g. Hull (2000), p. 540.
[58] Through ICAP via Bloomberg.

terms. Hence, there are many things that should be taken into account when calibrating the Hull-White parameters from the cap and/or floor prices on the market. The quoted implied volatilities as of November 21, 2005 are listed in Table 3.3.

Maturity (years)	EUR	DKK	DKK/EUR mark-up	EUR-DKK
1	17.62 %	19.90 %	1.13	2.3 %-points
2	20.36 %	22.70 %	1.11	2.3 %-points
3	21.77 %	23.40 %	1.07	1.6 %-points
4	21.44 %	23.50 %	1.10	2.1 %-points
5	21.22 %	23.30 %	1.10	2.1 %-points
6	20.92 %	22.90 %	1.09	2.0 %-points
7	20.59 %	22.50 %	1.09	1.9 %-points
8	20.19 %	22.10 %	1.09	1.9 %-points
9	19.84 %	21.70 %	1.09	1.9 %-points
10	19.50 %	21.40 %	1.10	1.9 %-points
15	17.99 %	N/A	N/A	N/A
20	17.01 %	N/A	N/A	N/A

Source: ICAP via Bloomberg

Table 3.3: Quoted cap (flat) volatilities as of November 21, 2005

Euro cap volatilities and DKK cap volatilities with the two different extrapolation assumptions are also shown in Figure 3.5. Notice the peculiar hump shape of the flat volatilities. This is a commonly observed phenomenon, but has actually been somewhat a puzzle. Hull (2000) suggests that the existence of the hump shape may be due to the *sources* of uncertainty distributed along the maturity spectrum. The short rates are to a large extent controlled by central banks, and have, therefore, limited volatility. For the long rates, the mean reversion property of the interest rate evolution process causes volatilities to decline. However, medium term interest rates are to a large extent determined by supply and demand in fixed income markets. The hypothesis is that investors tend to overreact to market movements, causing volatilities in this spectrum to be relatively large.

The calibration could in principle be conducted in an Excel spreadsheet very much like the way we set up the Nelson-Siegel yield curve estimation routine. However, the gain of setting up the calibration program ourselves would be further time consuming, and the gain of carrying out this task does not outweigh the time costs in our opinion. We therefore apply a professional piece of software to carry out the calibration, namely

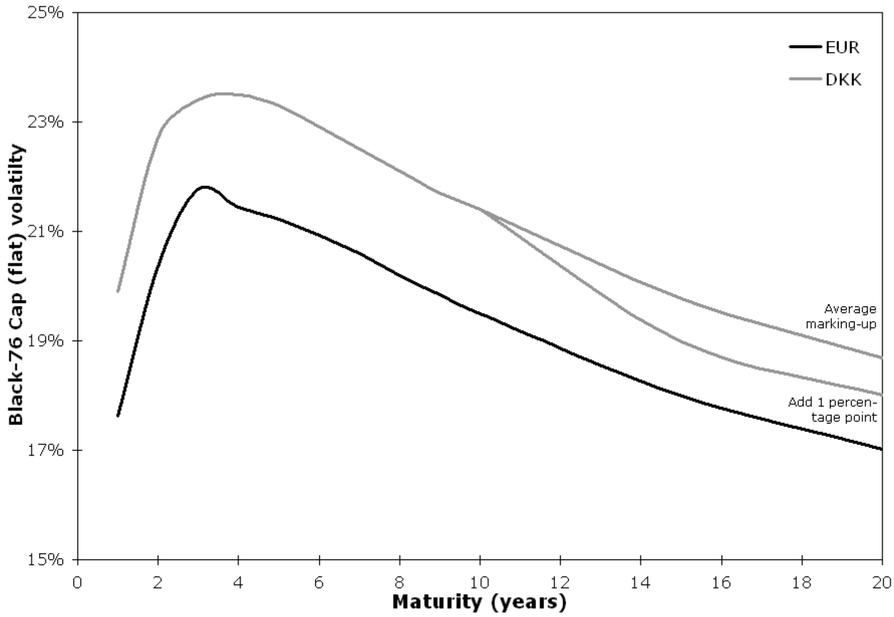

Figure 3.5: Cap volatilities as of November 21, 2005

the FinE Function Library.[59] The calibrated parameters for three different assumptions of how to deal with caps of longer maturities are shown in Table 3.4.

Assumption for maturities of 15 and 20 years	$\hat{a}$	$\hat{\sigma}$
DKK Cap vol = EUR Cap vol + 100 bp	0.02906	0.8568%
DKK Cap vol = EUR Cap vol · Average mark-up (1.09)	0.01570	0.8312%
DKK Cap vols omitted	0.00010	0.7977%

Source: Own calculations conducted in FinE Function Library

Table 3.4: Hull-White calibrated parameters

From Table 3.4, it is evident that the volatility parameter σ seems to be rather stable, while the mean-reversion parameter a on the other hand seems to be rather unstable. It is a well-known problem with the Hull-White model that the mean reversion parameter is fairly unstable. Some practitioners go as far as to suggest to fix the parameter a at a reasonable level, and only estimate the volatility parameter σ. We do not follow this advise, but proceed with the solution obtained following industry practice, namely adding 1 percentage point to the Euro cap volatilities for maturities of 15 and 20 years. From Table 3.4, we see that this gives rise to parameter estimates of $\hat{a} = 0.02906$ and $\hat{\sigma} = 0.8568\%$. These are the parameter estimates that we will use when applying the Hull-White model in the subsequent section.

[59] www.fineanalytics.com. The software is kindly made available by FinE Analytics.

3.4 Implementing Hull-White

Now that we have both estimated a yield curve and calibrated the Hull-White model, we turn towards the issue of how to implement the model. In specific, what we would like to do, is to use the Hull-White model to generate a range of possible interest rates in future periods. These future interest rates can be used to estimate the likelihood of the embedded prepayment options being exercised in future periods. In other words, the cash flow of a callable mortgage bond is uncertain, and we therefore need a model for the future interest rates such that we can estimate how large a fraction that will be prepaid in future periods. We return to prepayment issues in section 4, but before we do that, we go through the implementation of the Hull-White model in this section.

There are multiple ways to apply the Hull-White model. Among the most used are Monte Carlo simulation and interest rate tree building. Monte Carlo simulation is a good method since it is a very general procedure that is very suitable for valuing also path-dependent products. Monte Carlo simulation requires considerable computational power, and it becomes more and more applied as technology advances. However, for expositional purposes, we choose to do interest rate tree building. This method provides us with a good insight into how the model works, and how the parameters influence the evolution of the short interest rate.

An interest rate tree is a way to represent the stochastic process for the instantaneous short-term interest rate in discrete time. The interest rate trees often take the trinomial form. This simply means that from every node in the tree, it is assumed that the interest rate in the next period can take on one of three possible values. The branching inside the tree, however, may vary from node to node. The three possible branching methods in the trinomial tree are shown in Figure 3.6.

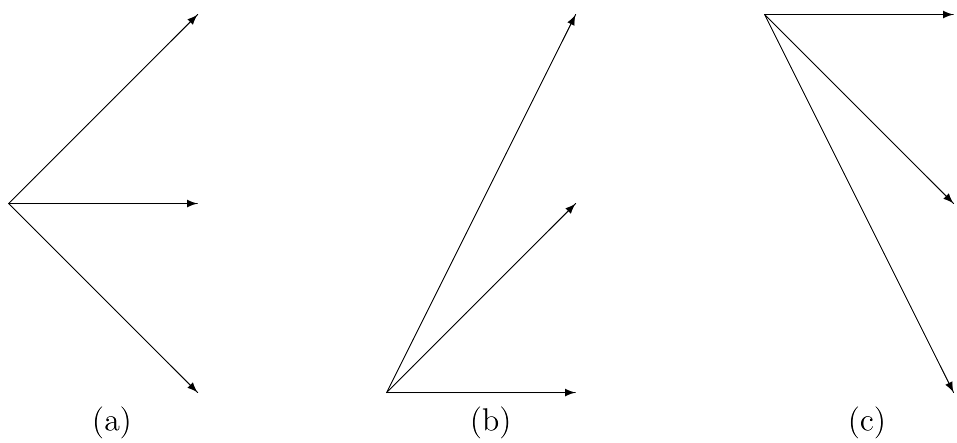

Figure 3.6: The branching methods in a trinomial tree

Hull & White (1994) pioneered the use of trinomial interest rate trees by making a discrete representation of stochastic term structure models. The trinomial tree obviously distinguishes itself from a binomial tree by providing an extra degree of freedom. Hull (2003) states that this enables the interest rate tree to represent e.g. mean-reversion more easily than with a binomial tree.[60] Previously, it has been common to use the simpler binomial representation for term structure models, for instance the Black, Derman & Toy (1990) model or the Black & Karasinski (1991) model.

The interest rate tree building logically consists of two parts:[61]

- Creation of an interest rate tree for an auxiliary variable R^* that is initially zero.
- Transformation of the interest rate tree for R^* into a tree for the short-term interest rate R.

It is natural to make the assumption that the discrete time short interest rate follows the same stochastic process as the instantaneous interest rate, and this is exactly what we will do. In the Hull-White model, the instantaneous short rate r follows the process

$$dr = [\theta(t) - ar]dt + \sigma dW^Q \tag{3.50}$$

Hence, we will now assume that the discrete time (Δt) interest rate R follows the same stochastic process:

$$dR = [\theta(t) - aR]dt + \sigma dW^Q \tag{3.51}$$

Note that in the limit where $\Delta t \to 0$, the two processes converge, so the assumption seems fair. We now define a new variable R^* by setting $\theta(t) = 0$. R^* has the property of being zero initially and it develops according to:

$$dR^* = -aR^*dt + \sigma dW^Q \tag{3.52}$$

We will later need the first and second moments of the distribution of the discrete-time change variable $[R^*(t + \Delta t) - R^*(t)]$. The distribution can be shown to be[62]

$$[R^*(t + \Delta t) - R^*(t)] \sim N(-aR^*(t)\Delta t, \sigma^2 \Delta t) \tag{3.53}$$

[60] Hull (2003), p. 551.
[61] For thorough references to the exposition in the following, see Hull (2003) or Hull & White (1996).
[62] Hull (2000) p. 581.

We start by going through the steps needed to create an interest rate tree for the variable R^*. What we need to do first, is to determine the overall shape of the tree. We need to decide the length of the constant[63] time steps Δt on the tree. We choose a time step of three months, i.e. $\Delta t = 0.25$. This is a natural choice, since the ultimate goal of the exercise is to price mortgage bonds, which often have quarterly payment dates and quarterly Bermudan-style prepayment option exercise dates. So what we really need, is to know the interest rate on the dates of possible exercises, which occur once every quarter. This way of modelling mortgage bonds can be problematic, since, even though the prepayment option can only be exercised at a payment date, most mortgage banks offer borrowers an opportunity to take on a new loan between payment dates at the present price, e.g. when prepaying their existing loan. So, borrowers actually *do* have the opportunity to act on beneficial interest rate movements in between two interest payment dates. However, we stick to the time span of three months between nodes on the tree, knowing that this is indeed an approximation.

Next, we need to determine the difference in the interest rate ΔR^* between two vertically adjacent nodes on the tree. Hull & White (1994) argue that

$$\Delta R^* = \sigma \cdot \sqrt{3 \cdot \Delta t} \tag{3.54}$$

is a good choice from the standpoint of error minimization. In our case this means that $\Delta R^* = 0.8658\% \cdot \sqrt{3 \cdot 0.25} = 0.7498\%$. The next thing we need to decide is which branching method to use in the tree. In order to do this, we introduce some notation of the nodes in the tree. A node is identified by a set of integer coordinates (i, j), where $t = i \cdot \Delta t$ and $R^* = j \cdot \Delta R^*$. Hence, this corresponds to defining a coordinate system with its starting point in the initial (time 0) node. $i \geq 0$ is the horizontal (time) distance from the initial node, while $j \in \{j_{\min}, \ldots, -2, -1, 0, 1, 2, \ldots, j_{\max}\}$ is the vertical distance from the initial node. Often, bounds are imposed on $j_{\min}$ and $j_{\max}$ in order to ensure that the probabilities in the tree are always non-negative. This means that when $j = j_{\min}$, branching type (c) from Figure 3.6 is used, when $j = j_{\max}$, branching type (b) is used, and finally when $j_{\min} < j < j_{\max}$, the branching type (a) is used. Hull & White (1994) claim that the most efficient way to calculate $j_{\min}$ and $j_{\max}$, which at the same time ensures that the probabilities are non-negative, is to set $j_{\max} = \text{Integer}\left(\frac{0.184}{a \cdot \Delta t}\right) + 1$ and $j_{\min} = -j_{\max}$.

Now we have almost got everything we need in order to create the interest rate tree

[63]The tree-building procedure can be extended to accommodate interest rate trees with non-constant time steps, but we will refrain from showing it here.

for R^*. The only thing missing is to derive the (martingale) probabilities (of each of the three possible outcomes in the next period) in all nodes. Here we show the derivation for branching type (a). To calculate these probabilities, we use the first and second moments of the distribution of $R^*(t+\Delta t) - R^*(t)$ (in the tree) and we match these moments with the probabilities in the tree. This gives us two equations with three unknowns, so we still need one more equation in order to determine the probabilities. The last equation is just that the sum of the probabilities must sum to one. Hence, we have the following three equations in three unknowns; the three probabilities, p_u, p_m and p_d:

$$p_u \Delta R^* - p_d \Delta R^* = -aj\Delta R^* \Delta t \tag{3.55}$$

$$p_u(\Delta R^*)^2 + p_d(\Delta R^*)^2 = \sigma^2 \Delta t + a^2 j^2 (\Delta R^*)^2 (\Delta t)^2 \tag{3.56}$$

$$p_u + p_m + p_d = 1 \tag{3.57}$$

In the appendix A.1 these three equations are solved, and the results are shown to be:

$$p_u^{(a)} = \frac{1}{6} + \frac{a^2 j^2 (\Delta t)^2 - aj\Delta t}{2} \tag{3.58}$$

$$p_m^{(a)} = \frac{2}{3} - a^2 j^2 (\Delta t)^2 \tag{3.59}$$

$$p_d^{(a)} = \frac{1}{6} + \frac{a^2 j^2 (\Delta t)^2 + aj\Delta t}{2} \tag{3.60}$$

When using branching type (b), the probabilities can be shown to be[64]

$$p_u^{(b)} = \frac{1}{6} + \frac{a^2 j^2 (\Delta t)^2 + aj\Delta t}{2}$$

$$p_m^{(b)} = -\frac{1}{3} - a^2 j^2 (\Delta t)^2 - 2aj\Delta t$$

$$p_d^{(b)} = \frac{7}{6} + \frac{a^2 j^2 (\Delta t)^2 + 3aj\Delta t}{2}$$

Finally, when using branching type (c), the probabilities can be shown to be

$$p_u^{(c)} = \frac{7}{6} + \frac{a^2 j^2 (\Delta t)^2 - 3aj\Delta t}{2}$$

$$p_m^{(c)} = -\frac{1}{3} - a^2 j^2 (\Delta t)^2 + 2aj\Delta t$$

$$p_d^{(c)} = \frac{1}{6} + \frac{a^2 j^2 (\Delta t)^2 - aj\Delta t}{2}$$

[64] Hull & White (1994), p. 11.

To show how the method works, we now calculate an interest rate tree for the first two years, i.e. the first eight quarters. Since, in this case, $j_{\max} = 26$, we use branching type (a) and the corresponding probabilities (3.58) – (3.60) in all nodes, since the maximum reachable j is 8. The interest rate tree for R^* is shown in Figure 3.7. This interest rate tree is in itself of little interest, but it can already at this stage, be used to see the maximal changes in the interest rate from the starting point to a given time step according to the model.

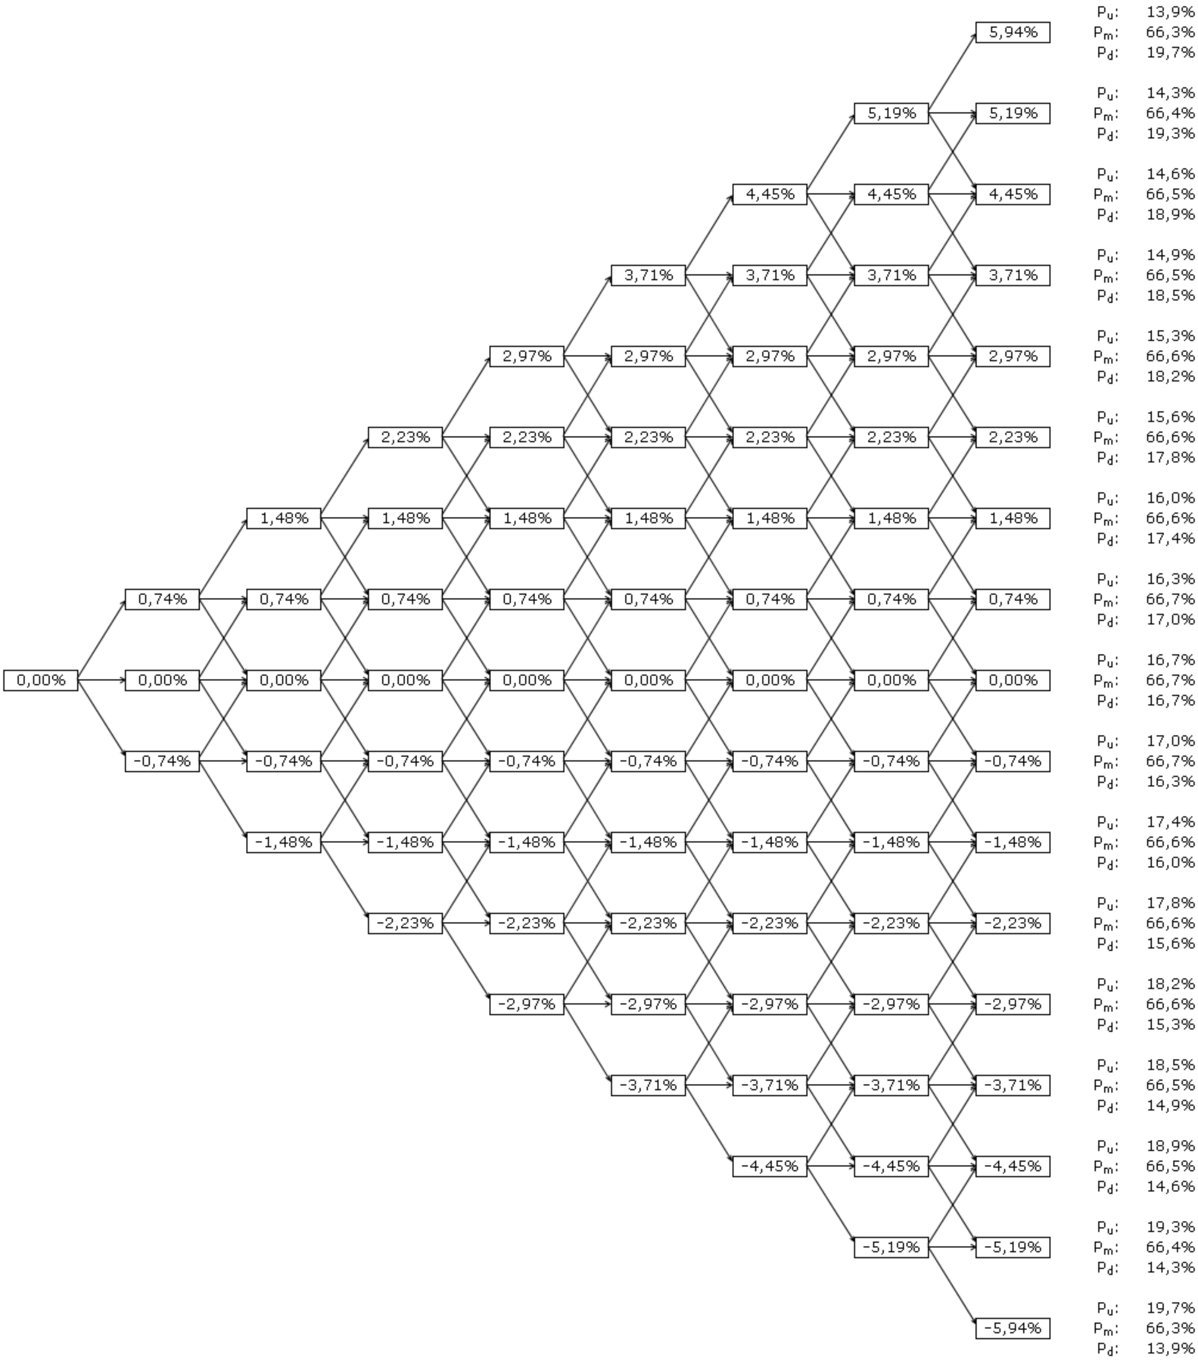

Figure 3.7: Interest rate tree for R^*

Now that we have shown how to make an interest rate tree for R^*, we will turn the view towards how to transform the interest tree for R^* into an interest rate tree for R.

To begin with, we define a new variable $\alpha(t)$ by

$$\alpha(t) = R(t) - R^*(t) \tag{3.61}$$

Note that if we can calculate $\alpha(t)$, we would immediately also know $R(t)$, which is the aim of the entire exercise, and note furthermore that $E[\alpha(t)] = E[R(t)]$, since $E[R^*(t)] = 0 \,\forall\, t$. For notational convenience we denote $\alpha_i \equiv \alpha(i\Delta t)$. We need yet another variable $\mathcal{G}_{i,j}$, which is the present value of a security that will give a payoff of 1 if the node (i, j) in the tree is reached, and 0 otherwise. We use these auxiliary variables ($\mathcal{G}_{i,j}$'s) to calculate the α_i's, and hence create the interest rate tree for R by adding R^* to α_i.

The overall idea in the calculations, which will be conducted using forward induction, is to match the value of a zero coupon bond with the value of the full collection of $\mathcal{G}_{i,j}$'s with the same maturity, such that this collection also exactly gives a payoff of 1 with certainty. In other words, the idea is to match the value of zero coupon bond with the value of a synthetic portfolio of other securities (the $\mathcal{G}_{i,j}$'s) that in total has a payoff profile exactly equal to that of a zero coupon bond.

The value of a zero coupon bond with principal 1 maturing at time $(i+1)\Delta t$ is

$$P_{i+1} = e^{-\tilde{r}_{i+1}(i+1)\Delta t} \tag{3.62}$$

where $\tilde{r}_{i+1}$ is the term $(i+1)\Delta t$ spot interest rate as measured by the initial yield curve, which we have already derived in section 3.1. We here make use of the derived initial yield curve. The idea is to match the value of this zero coupon bond with the expected value of the synthetic portfolio described above:

$$\begin{aligned} P_{i+1} = e^{-\tilde{r}_{i+1}(i+1)\Delta t} &= \sum_{j=-n_i}^{n_i} \mathcal{G}_{i,j} \cdot e^{-(\alpha_i + j\Delta R^*)\Delta t} \\ &= e^{-\alpha_i \Delta t} \cdot \sum_{j=-n_i}^{n_i} \mathcal{G}_{i,j} \cdot e^{-j\Delta R^* \Delta t} \end{aligned}$$

Taking logs and rearranging yields

$$\begin{aligned} -\tilde{r}_{i+1}(i+1)\Delta t &= -\alpha_i \Delta t \cdot \log\left[\sum_{j=-n_i}^{n_i} \mathcal{G}_{i,j} \cdot e^{-j\Delta R^* \Delta t}\right] \quad \Leftrightarrow \\ \alpha_i &= \frac{\log\left[\sum_{j=-n_i}^{n_i} \mathcal{G}_{i,j} \cdot e^{-j\Delta R^* \Delta t}\right]}{\Delta t} + \tilde{r}_{i+1}(i+1) \end{aligned} \tag{3.63}$$

where n_i is the number of nodes on each side of the central node in stage i. When α_i is determined, the $\mathcal{G}_{i+1,j}$'s can be determined through

$$\mathcal{G}_{i+1,j} = \sum_k \mathcal{G}_{i,k} \cdot q(k,j) \cdot e^{-(\alpha_i + k\Delta R^*)\Delta t} \tag{3.64}$$

where $q(k,j)$ is the probability of moving from node (i,k) to node $(i+1,j)$. The summation is done for all nodes in the previous stage, of which some may have an attached probability of zero. By the use of (3.63) and (3.64), we can iteratively calculate α_i's and $\mathcal{G}_{i,j}$'s through the tree using a forward induction principle. We now only need one more thing, namely a starting condition. This is obviously $\mathcal{G}_{0,0} = 1$. The value at time 0 for a bond that pays off exactly 1 at time 0 is of course equal to its payoff, 1. With $\mathcal{G}_{0,0}$ at hand, we can calculate α_0 as $\alpha_0 = \frac{\log(1 \cdot e^0)}{0.25} + \tilde{r}_1 = \tilde{r}_1$. The calculated tree for the values of $\mathcal{G}$ are shown in Figure 3.8 and the calculated values of α_i are shown in Table 3.5, where we apply the estimated (Nelson-Siegel) yield curve, which is shown in Figure 3.3 on page 33.

i	0	1	2	3	4	5	6	7	8
t	0.00	0.25	0.50	0.75	1.00	1.25	1.50	1.75	2.00
$\tilde{r}_i$	2.122%	2.174%	2.224%	2.274%	2.323%	2.370%	2.417%	2.462%	2.507%
α_i	2.174%	2.275%	2.374%	2.470%	2.564%	2.655%	2.743%	2.830%	2.913%

Table 3.5: Zero coupon interest rates and the auxiliary variable α_i

Since we now know α_i for all i, we can proceed to calculate the discrete time interest tree for the short rate R. We can already infer the expected development in the discrete time short rate R from (3.61), since the expected development in the short interest rate corresponds to the development in α_i as shown in Table 3.5. The interest rate tree for R is shown in Figure 3.9. From the interest rate tree, it is apparent that the possibility of negative interest rates is not just an academic issue, but indeed, it is evident that a significant share of the nodes in the interest tree does indeed have negative interest rates. This is, of course, a serious problem, but when choosing to implement the Hull-White

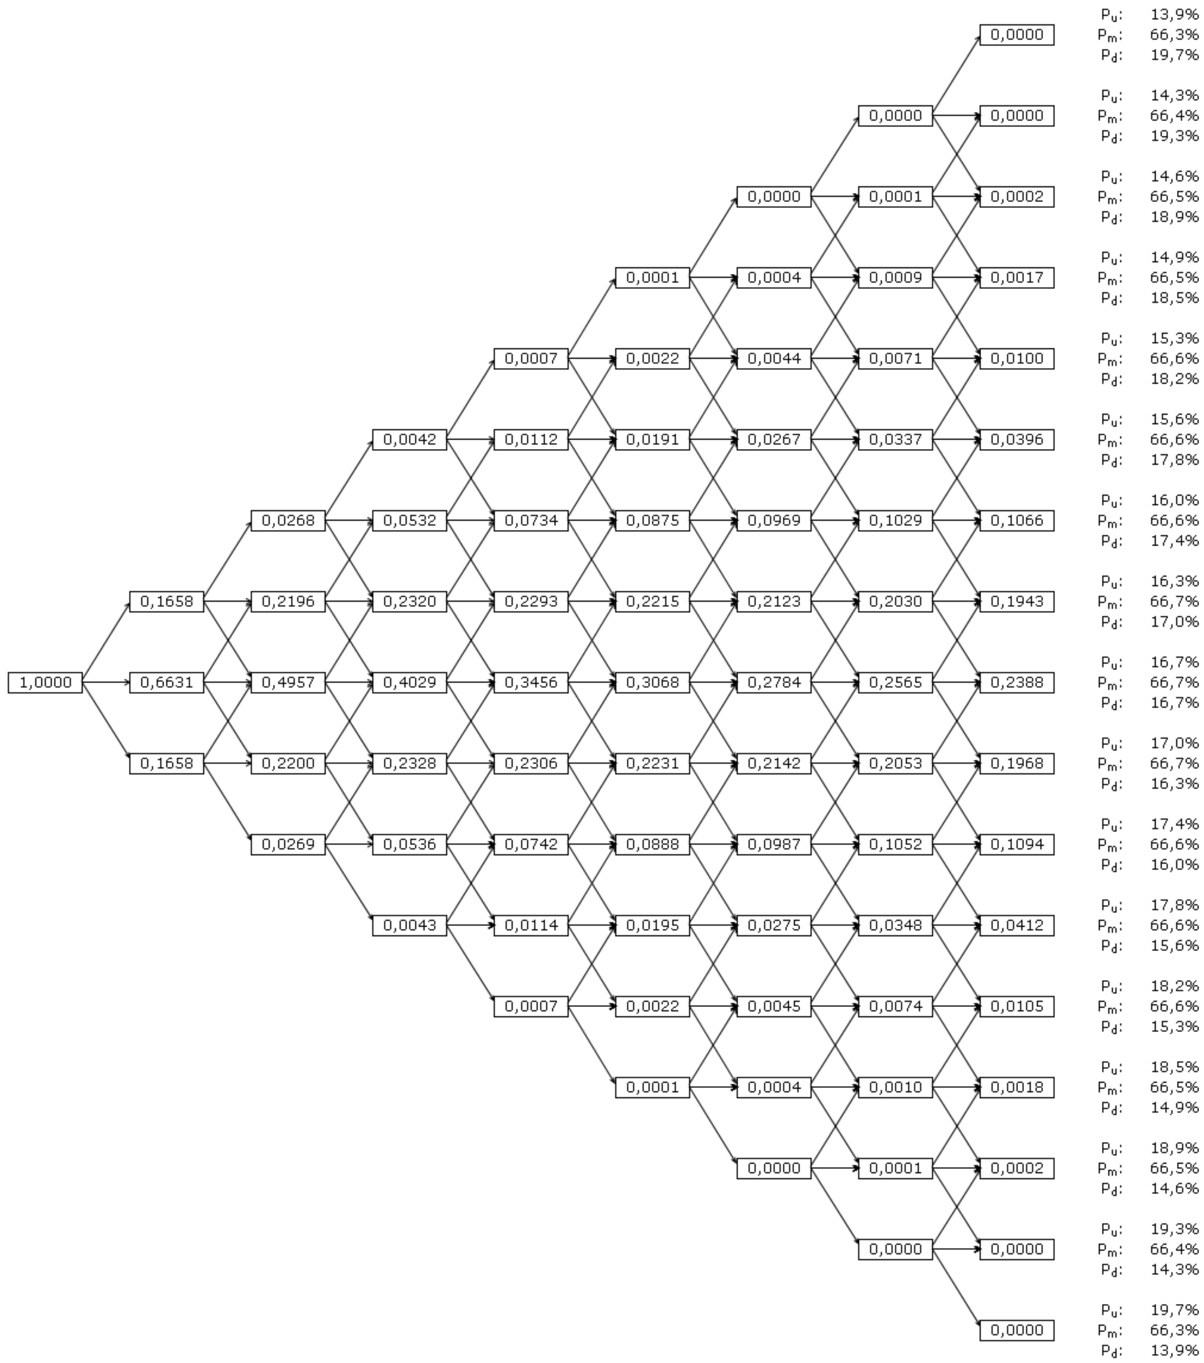

Figure 3.8: Tree for the auxiliary variable $\mathcal{G}$

model, this is something that one will have to accept. The reason why the problem of negative interest rates in the interest rate tree is as pronounced as it is in this case, is obviously closely related to a historically relatively low level of interest rates at the time of estimation. Furthermore, a high σ and a low a will also contribute to a higher likelihood of negative interest rates.

Hence, Figure 3.9 measures an interest rate tree for the short-term interest rate, which

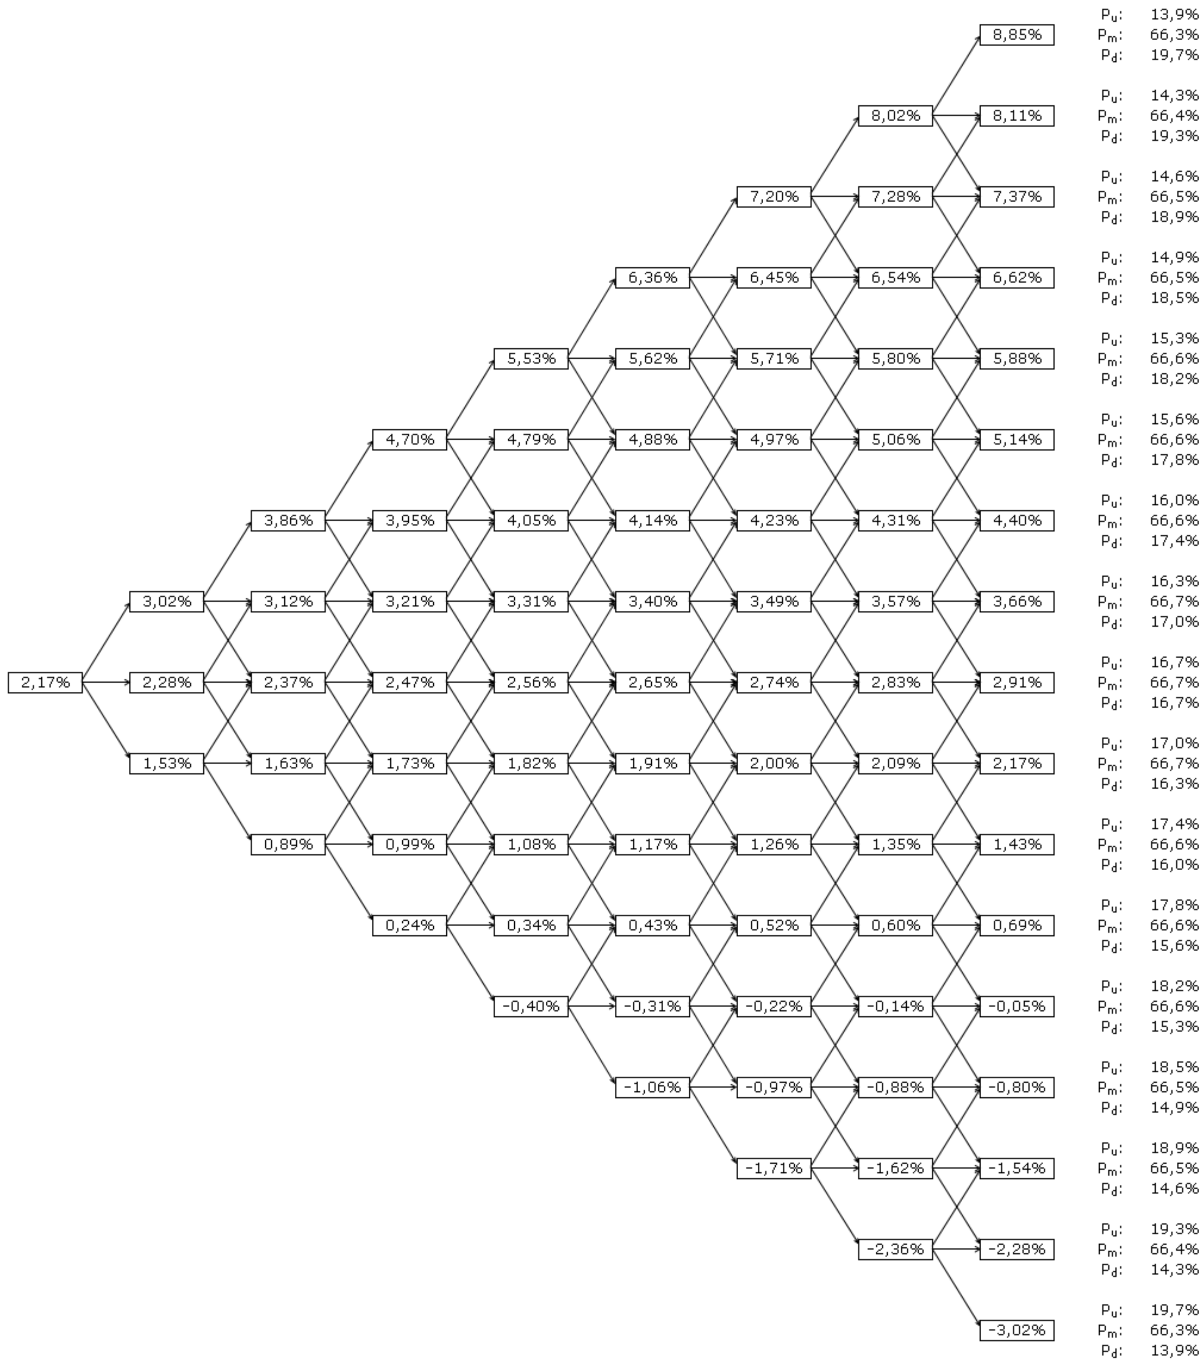

Figure 3.9: Interest rate tree the short term interest rate R

is the concrete outcome of the Hull-White model. The basic ingredients to create this tree are an initial yield curve, which is derived on basis of a selected sample of Danish mortgage bonds, and cap volatilities, which are used for estimating the mean-reversion parameter and the volatility parameter in the Hull-White model. This completes the first main part of setting up a mortgage bond pricing model. We now have an idea of the future interest rates based on a term structure model. In the next sections, we build the second part of the mortgage bond pricing model; prepayment modelling. We look closely into this issue in the coming sections.

4 Prepayment Behavior

In the two preceding sections, we have developed a pricing model, which enables us to price any cash flow. Actually, to price an asset with a deterministic cash flow – for instance a non-callable mortgage bond, we only need a relevant yield curve. The callability of a traditional Danish mortgage bond is, as noted in the introduction, what complicates things considerably. In other words, it is the uncertainty of the cash flow of a callable bond that makes it particularly difficult to price. Hence, we need to develop an extension to the existing pricing model in order to price callable mortgage bonds. This extension is a prepayment model, for which an important prerequisite is a model for the evolution of the term structure of interest rates. The reason why we need to model the evolution of the term structure, is to obtain the value of the prepayment option in the future, since this enables us to estimate the size of prepayments in future periods. The issue of modelling prepayments have only been treated to a limited extent in a Danish context, since the prepayment models are usually developed by e.g. commercial banks. Hence, the academic literature on prepayments with special emphasis on the Danish case is relatively scarce.

Once it is noted that a traditional Danish mortgage bond consists of a non-callable bond and a sold call option, it is a natural suggestion to price these two assets separately and calculate the total value of the callable bond as the value of a non-callable bond with similar properties subtracted the value of the call option, cf. equation (1.1). Christensen (2005) shows how to value a callable Danish mortgage bond using a preliminary approach, applying known option pricing formulas; in particular a modified version of the Black & Scholes (1973) model and a binomial model, respectively. However, the results are not satisfactory. Christensen (2005) concludes that the Black-Scholes model for valuing the prepayment option can provide an approximate suggestion of the value of the prepayment option, but it is not suitable to make a reasonable model for valuing callable mortgage bonds.

Instead of modelling the price of a callable mortgage bond by valuing a non-callable bond with similar properties and a call option on the bond analytically and separately, the usual way to value the callable mortgage bond is to use some sort of a prepayment model.

The main objective of a prepayment model is not to predict future prepayments, but to establish a connection between projected mortgage rates and projected prepayments. By modelling mortgagors' prepayment behavior for a given mortgage rate, we can, by projecting future mortgage rates using our term structure model, also project future prepayments. Combining the term structure model with a prepayment model, we can obtain the fair value of a callable mortgage bond. We will briefly discuss how to combine the term structure model and the prepayment model in section 6. Before we get to that, we focus on the properties of prepayment behavior in this section, and the set-up of a model for prepayments in section 5.

4.1 Prepayments in General

Before we address the issue of modelling prepayments, we need to establish more precisely, how prepayments are measured. Prepayments are usually measured as the *Conditional Prepayment Rate* (CPR). CPR is the percentage part of the total outstanding amount at a given point in time that is prepaid. So if the CPR is 25%, this simply means that 25% of the notional amount in a bond series is being prepaid in that period. From the investors' point of view, this means that 25% of the holding is redeemed at par, leaving the investor with an investment of only 75% of the nominal amount before the extraordinary redemptions.[65] When setting up prepayment models, it is usually done by the use of CPR.

To begin with, a natural question to ask would be, whether the issue of prepayments is really a significant issue. Is it really worth all the trouble going through advanced pricing models? If prepayments are a phenomenon of insignificant importance, it does not seem logic to spend a lot of effort on the explanation of its size. However, Figure 4.1 clearly shows that modelling prepayments is indeed necessary. This is due to two facts; the large size of prepayments, sometimes more than DKK 100 bn. in just one term, and the variance of the prepayment extent. The cash flow from a mortgage bond is thus greatly influenced by prepayments, and it is therefore very important to include in a valuation model.

Not surprisingly, the vast majority of the literature on prepayments is dealing with the American mortgage bond market. Furthermore, much of this literature additionally contains a large share of papers on proprietary models from investment banks promoted by the economic incentive that research in this field entails. We focus on the Danish set-up, but we will throughout this and the next section make references to and use papers also treating the American market.

As we mentioned in the beginning, the American and the Danish mortgage markets have many similarities, but there are some very important differences when it comes to

[65] This means that if the CPR is constantly 25% in four consecutive quarters, this leaves the investor with an investment of just $(1 - 0.25)^4 = 32\%$ after a year, ignoring ordinary redemptions. The rest of the investment has been redeemed at par.

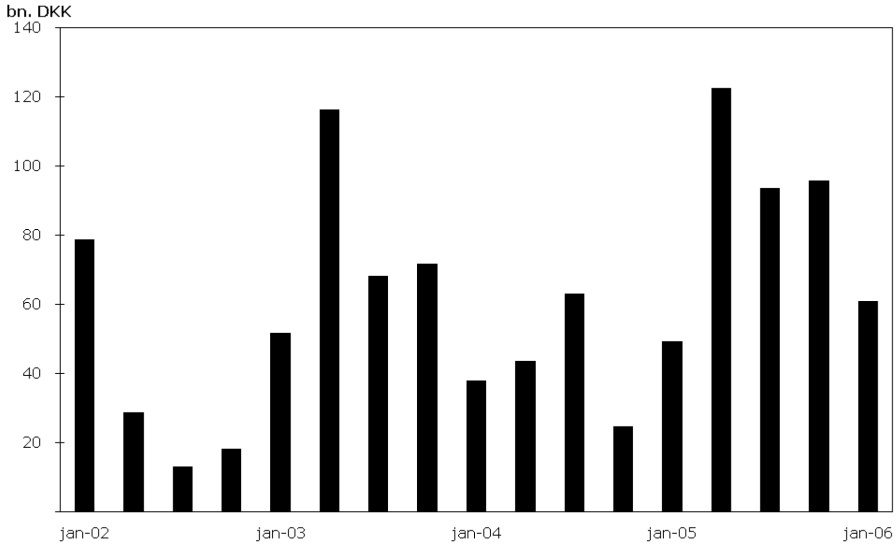

Source: Danske Research

Figure 4.1: Total prepayments on Danish mortgage bonds

the prepayment set-up. An American debtor has a standard call option on his mortgage, whereas a Danish debtor has both a call option *and* a delivery option. A Danish debtor can thus choose whether to call the option at the strike price (at par for a regular callable bond) or to buy an equivalent notional amount in the market (at market price) and then cancel the debt with the mortgage bank. This implies that if a Danish mortgagor wishes to cancel his loan, the cost of redeeming the loan is market value capped at nominal principal. Hence, the added delivery option effectively means that almost no prepayments occur in the Danish market as long as the bond price is below par. As the delivery option has a non-negative value, the existence of it decreases the value of a Danish callable bond compared to a American callable bond. The Danish mortgagor thus compensates the investor by paying a higher yield. In the case where the mortgagor chooses to exercise the delivery option, the investor does not incur a loss as the bond is purchased at market price. However, he incurs a loss compared to the American set-up where the mortgagor would have to prepay the loan at par and thereby pay a premium compared to the market value of the loan.

Later on, in section 5.4, we look in detail on the timing of prepayments. However, to ease the presentation and understanding of the process of prepayments in the Danish case, we briefly go through the time line of a typical Danish mortgage bond. A traditional Danish mortgage bond has four yearly terms, at the beginning of January, April, July and October. The mortgagor must announce that he wishes to exercise his prepayment option no later than two months before the relevant due date, which means that the closing dates of exercise of the prepayment options are at the *end* of January, April, July and October. Therefore, if a mortgagor wishes to prepay at the April term he must announce it before January 31st. The call option is thus a so-called Bermuda option[66] as the option can only be exercised at predetermined dates throughout the life of the option.[67]

In the following, we initially analyze the decision that a rational debtor faces concerning the optimal strategy for his prepayment option. The section on rational prepayment behavior is followed by a presentation of a list of important drivers of prepayments. In section 5, we proceed to present two prepayment models, one targeted at the American case, and one targeted at the Danish case. This leads us to the set-up of our own prepayment model in section 5.3, where we will make use for the observations regarding prepayment behavior from the present section.

4.2 Rational Prepayment Behavior

As hinted by the name, a rational prepayment model assumes that the mortgagor acts rationally in his prepayment decision. The literature on rational prepayment models is truly vast, and we therefore merely aim at presenting the basic idea of this model class to facilitate the reader's understanding of the incentives behind prepayments.

Most rational prepayment models do not have a closed-form solution and will therefore need to be solved numerically. We do not carry out numerical solutions, as we present the rational prepayment model mainly to facilitate the reader's understanding of the complexity of the valuation of the prepayment option.

In Brennan & Schwartz (1977), the prepayment decision for an American call option on a zero coupon bond is modelled in a continuous framework.[68] They apply the perspective that a mortgagor seeks to minimize the value of his liabilities as a necessary condition for maximizing net present value. It is assumed in the paper that arbitrage opportunities do not exist and that markets are frictionless. The authors use the term structure equation, which we have derived in Result 2.1, for a non-callable bond and, subsequently, intuitively derives the optimal call strategy. Markets are frictionless, which implies that the mortgagor should call the loan whenever the bond price equals the strike price, which

[66] See Hull (2000), chapter 18 for a description of Bermuda options.

[67] Actually, the rules are a little bit more complicated, so in fact the prepayment option embedded in a mortgage loan is only Bermudan-*style*, since the mortgagor can actually make a so-called immediate par redemption (Danish: "pari-straks"), but since this business is carried out with the mortgage bank as the counterpart and not the investor, we can say that the prepayment options embedded in traditional Danish mortgage bonds are Bermuda-options.

[68] Dunn & McConnell (1981) extend the Brennan-Schwartz model by among other things using a amortizing bond.

equals par. Calling it below par will obviously be suboptimal as the value of the debt is lower than the cost of calling the loan. Calling it at a price above par is also suboptimal as the mortgagor could have decreased the value of the debt by calling it at an earlier point in time. Formally, the model thus implies the following prepayment behavior

$$CPR(r, L, t, T) = \begin{cases} 1 & \text{for} \quad F(r, L, t, T) > L \\ 0 & \text{otherwise} \end{cases} \tag{4.1}$$

where $F(\cdot)$ denotes the bond price and L denotes the notional amount. Hence, the model predicts prepayment behavior that is solely dictated by the bond price. This implies that there exists a critical yield, r^*, defined such that $F(r^*, \cdot) = L$ at which the loan is prepaid. The CPR can thus be rewritten in yield terms as $CPR(r, t) = 1_{\{r_t < r^*\}}$. Having prepayment behavior, which is dictated alone by the bond price, implies that mortgagors within the same bond series all prepay at the same time, independently of individual loan characteristics. Hence, when pricing a callable zero coupon bond paying $\$L$ at time T, Brennan & Schwartz (1977) add the following boundary condition

$$F(r, L, t, T) \leq L \quad , \quad \forall \;\; t < T \tag{4.2}$$

to the term structure equation.

The dynamics of the Brennan-Schwartz model can be described as follows. When the bond price supersedes par, all loans are prepaid and the bond series close. These loans are subsequently refinanced with a loan with an infinitesimally lower net present value and coupon rate and is priced at par. If the rate decreases further, the afore-mentioned routine would be carried out again. In Figure 1.1, we saw that market prices of callable mortgage bonds are not capped at par as predicted by the Brennan-Schwartz model. However, for a very parsimonious model it captures the gist of it (though intuitively obvious), which is that the price of a callable bond has limited upside potential due to the call option.

As mentioned, there does not exist a closed-form solution to the Brennan-Schwartz model and it is therefore solved numerically using e.g. estimation via a pricing tree as the one we have constructed in section 3.4.[69] By merely adding (4.2) to the pricing tree, which caps the bond price at par, we obtain the Brennan-Schwartz bond price. This bond price is, of course, lower than the non-callable bond price as the shorted call limits the upside potential of the bond price.

[69]One can also apply the finite difference method. We refer to Hull & White (1990*b*) for more on this method.

A much needed extension to the Brennan-Schwartz framework is the inclusion of transaction costs. Say, the mortgagor incurs a loan size dependent cost, $X(L)$, when prepaying his loan. We still assume that mortgagors minimize the present value of their debt and that the mortgagors instantaneously optimize the value of their loan. Therefore, we can write the value of the debt as the minimum of the debt value if the mortgagor prepays the loan and the debt value if the mortgagor does not exercise the option. Formally, this means that the value of the debt V_t at time t is given by

$$V_t = \min\{F(r, L, t, T); L + X(L)\} \tag{4.3}$$

where $F(\cdot)$ is the value of the existing debt given that the loan is not prepaid, L is the notional, and $X(L)$ is the cost of prepaying a loan of size L. When calculating V_t, a rational mortgagor is taking the entire term structure – according to the Hull-White model in our set-up – into account. If he chooses to prepay his loan, he must pay the remaining principal L and the prepayment costs, $X(\cdot)$.[70]

From (4.3), we infer that the optimizing mortgagor prepays his loan if $F(r, L, t, T) > L + X(L)$. This is completely analogue to the Brennan-Schwartz set-up. However, now we cannot define a global critical yield, since we now have two variables in play. Besides the refinancing interest rate, the loan size is also a determinant of prepayment behavior. If $X(L)$ is assumed to be non-decreasing in L, then we infer that the larger the size of the loan, the higher the critical yield, at which the loan is prepaid. Another way to put it is that the larger the size of the loan, the earlier the loan will be prepaid.

Thus, this extension to the Brennan-Schwartz model is a remedy to two of the shortcomings of this model. First, since the decision whether to prepay a loan or not is dependent on the loan size, this enables the model to incorporate running prepayments, since mortgagors will prepay at different times. Next, since mortgagors may not prepay their loans (because of the refinancing cost) even though the refinancing interest rate is lower than the coupon rate, the bond price is no longer capped at par, but at a level above par. This level is determined by the size and structure of the prepayment costs and the debtor distribution in loan sizes. Both the running prepayments and the existence of prices of callable mortgage bonds above par are properties that are observed in reality, and therefore it is expedient that the extended model can incorporate these features.

However, the extended model does not solve all problems. We note that this model set-up cannot provide us with continuous prepayments in a scenario with constant or

[70]For sake of simplicity, we ignore discreteness of interest payments.

increasing interest rates, though this is observed in reality. To cope with this, modelers have introduced a baseline prepayment level using a hazard function. This approach is most applicable to the American market due to the fact that mortgagors does not have a delivery option and that exogenous factors such as house sales, divorce etc. can lead to seemingly irrational exercise of the prepayment option. We refer to Stanton (1995) for a model using this approach.

We now move on to looking at various possible drivers of prepayments, leading to section 5, where we will look closely into another and much more applied class of prepayment models that build on the drivers that we present below.

4.3 Drivers of Prepayment Behavior

Rational behavior models can only provide a partial description of prepayment behavior. In this section we present a selection of the most important prepayment drivers, which are established relevant in the literature. This section serves to present the variables that will be included in the prepayment models in section 5. We look into the following prepayment drivers

- Economic gain
- Maturity and burn-out
- Loan size

4.3.1 Economic Gain

The single most important factor for triggering prepayments is inarguably the economic gain from exercising the prepayment option. Though we a priori do not believe that rational behavior provides a complete description, we believe that most enterprizes practice active debt management. Furthermore, most households must be expected to follow the advice of mortgage banks, which is based solely on the economic gain of prepaying. So, in total, it seems to be a fair assumption that the primary factor influencing prepayments is the economic gain.

We therefore wish to derive an estimate for the economic gain that can be realized by prepaying a mortgage loan. We have seen that a callable loan can be decomposed into a non-callable loan and a short call option. Hence, we wish to estimate the difference in the values between the current loan and the refinancing alternative.

How one defines the refinancing rate is of great importance for the derivation of the economic incentive. It is common practice to use the rate for a loan with similar characteristics, i.e. loan-type, maturity, coupon frequency etc. This is called the assumption of neutral behavior. The assumption of neutral behavior is clearly imperfect, and preferably, we would like to be able to analyze the pattern of mortgagors' choice of loan for refinancing. Such data could provide us with great insights into what drives the prepayment, but unfortunately it is not publicly available.

The interesting question is, what kind of behavior should be assumed instead of neutral behavior? If one compares different loan types, maturities etc., one inadvertently ends up comparing apples and oranges. One can easily measure the economic incentive from converting into a short-term non-callable loan, but the question is whether this is a relevant exercise. Mortgagors self-select themselves into different loan types with different corresponding risk profiles and can be expected only to migrate to a limited extent. We therefore argue that the simplifying assumption of neutral behavior may not be as restrictive as it immediately seems. Furthermore, it is a burdensome task to calculate an economic incentive for each of the refinancing alternatives a mortgagor faces, especially as the palette of loan alternatives expands cf. section 9.

The gain of prepayment arises from a difference in the coupon rate and the refinancing interest rate. If the difference is large, we expect prepayments to be higher, all other things equal. A natural suggestion would be to use $c - r$ as an indicator of the economic gain of prepayment. However, this measure has been criticized for being somewhat arbitrarily chosen.

Instead, Richard & Roll (1989) argue that it is more reasonable to calculate the present value of the annuity per unit of notional as an indicator. They do this by dividing the present value of an annuity with a constant quarterly payment of \$1

$$\mathrm{PV} = \frac{1 - (1+r)^{-T+t}}{r} \tag{4.4}$$

with the outstanding principal per quarterly payment of \$1

$$\mathrm{OP} = \frac{1 - (1+c)^{-T+t}}{c} \tag{4.5}$$

This yields the following expression

$$\frac{\mathrm{PV}}{\mathrm{OP}} = \frac{c}{r}\left[\frac{1 - (1+r)^{-T+t}}{1 - (1+c)^{-T+t}}\right] \tag{4.6}$$

The use of this indicator is intuitively appealing, since it compares the market value of the existing loan with the notional, which is the cost of prepayment.[71]

As Richard & Roll (1989), we use $\frac{c}{r}$ as an indicator for the economic gain, which can be seen to be a fairly good estimate of the economic gain from (4.6), provided that the term in the parenthesis is fairly constant. The higher the ratio, the higher the prepayment incentive. To control whether we can regard $\frac{c}{r}$ as being a prepayment driver, we investigate the co-movement between this fraction and observed CPR. To calculate the refinancing rates for the different loans can be a rather cumbersome assignment. From the perspective of the modeler, it poses a problem that the prepayment date can be chosen at the discretion of the mortgagor as we cannot determine the exact refinancing rate for all maturities. To cope with this, most practitioners use one of two approaches. The most straightforward approach is to use standard benchmark refinancing mortgage rates. These are available for maturities of 10, 20 and 30 years. Then the one closest to the time to maturity of the loan being prepaid is chosen as the refinancing interest rate.[72] Another way is to make use of a relevant yield curve, such that for a loan having 23 years to maturity, one merely uses the estimated rate for the 23 year interest rate on the yield curve.

In Figure 4.2, we use the first alternative and plot CPR for RD 6% 2032 together with the calculated expression of (4.6) and $\frac{c}{r}$. We use the 30 year mortgage bond benchmark yield lagged two months as refinancing rate. By lagging the refinancing rate, we incorporate that the announcement period is leading the mortgage term.

It can be seen that the two economic gain estimators are tracking the CPR to a reasonable degree. The simple measure, $\frac{c}{r}$, seems to be able to track CPR for RD 6% 2032 just as well as the more complex measure defined in (4.6).

To evaluate the attractiveness of prepayment, one can also apply a different gain measure. The taxation scheme favors interest payments over repayments, since interest payments are partially tax deductible.[73] The interest element of an annuity decreases with time for a given loan and consequently so does the tax shield from these interest payments. When prepaying a current loan and refinancing it with another loan, the interest element

[71]Notice that the use of this measure implies independence of the loan size.

[72]The segmentation does not need to be symmetric. For example, Madsen (2005) applies an upward skew segmentation.

[73]The tax deductibility has been gradually reduced over the years, and today interest payments are only deductible in the local taxes, which constitute between 50%-90% of a person's total taxes depending on personal income.

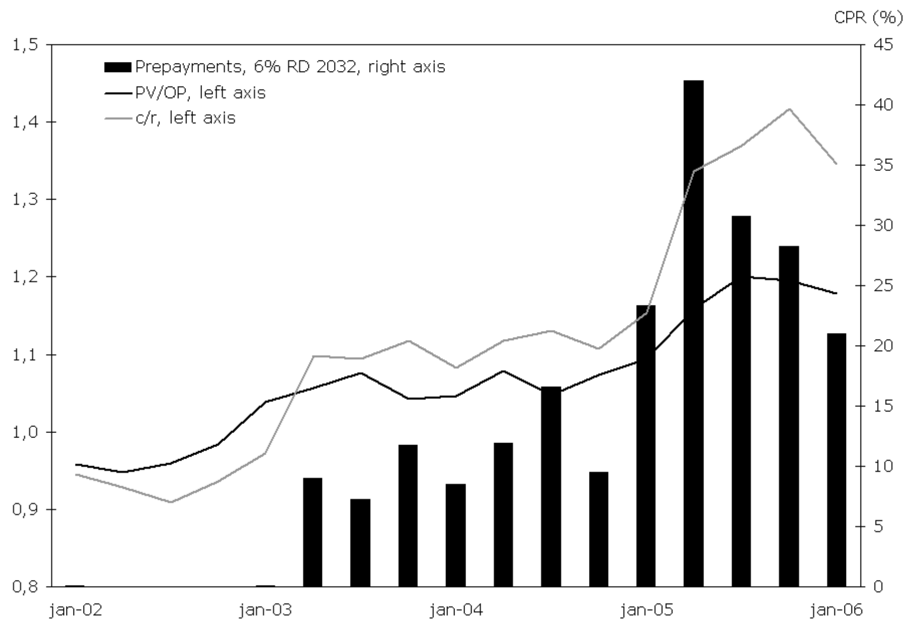

Source: Own calculations based on data from Danske Research and Danmarks Nationalbank

Figure 4.2: CPR and the economic gain

increases. We write the first year payment saving as

$$\text{FYP} = \frac{\text{FYP}^{\text{current}} - \text{FYP}^{\text{new}}}{\text{FYP}^{\text{new}}} \tag{4.7}$$

$\text{FYP}^{\text{current}}$ is not the actual first year of the current loan, but the first year of the remaining loan. One can easily imagine that household mortgagors are more prone to confuse first year payment and present value gains. However, most often these two will be intimately linked. The FYP-variable has historically been able to drive prepayments. Madsen (2005) conjectures that this tendency has decreased over recent years as mortgagors in general have been more driven by present value gains when exercising the prepayment option.

4.3.2 Maturity and burn-out

The time to maturity of a mortgage is often taken into account when modelling the prepayment extent. In the following we shed light on why this may be advantageous.

It does take quite an interest rate differential for it to be profitable to prepay a loan that lacks only a one-digit number of payments, due to existence of transaction costs. Hence, we expect loans with a short time to maturity to prepay only to a limited extent. Vice versa, loans with a long time to maturity must be expected to be prepaid to a much higher degree, simply because there are many payments left on the loan, and the gain of prepaying the loan is simply higher on average for loans with longer remaining maturity than for loans with shorter time to maturity. One can argue that this effect can be captured by an economic gain variable. This is a fair argument, and in fact one can say so about a lot of the variables that are often included in prepayment models. However, it is still fair to investigate whether these variables can contribute with a separate effect that is not captured by an economic gain variable. This will in particular be the case if the economic gain variable is an approximative variable. Since we use the approximative variable $\frac{c}{r}$, this underlines the relevance of a maturity variable in our case.

Connected to the issue of maturity is the concept of burn-out. The burn-out effect is the effect that series that have previously been through large waves of prepayments, tend to have lower prepayments than series in which this has not been the case. The burn-out effect is due to the fact that the mortgagors that are most eager to prepay their loans, and the mortgagors that are most observant to changing market conditions, have already prepaid their loans previously. Hence, the pool of mortgagors left in the series after waves of prepayments, are the mortgagors that are expectedly the most sluggish prepayers. To capture the effect of burn-out, one often uses the *pool factor*. The pool factor measures the ratio of the actual outstanding amount of a mortgage bond relative to the outstanding amount that would have been, had the series not been subject to prepayments:[74]

$$\text{Pool factor} = \frac{\text{Actual outstanding amount}}{\text{Outstanding amount in absence of any prepayments}} \tag{4.8}$$

Inclusion of the pool factor in the prepayment function actually complicates things considerably. This is due to the fact that the pool factor is path dependent. That is, the size of the pool factor in any given node in the interest tree not only depends on the actual node, but also on the path that is taken to reach that node. This makes the estimation of prepayments, and thereby the various possible cash flows, much more complicated. At (almost) every node in the interest tree, there are several paths leading to each node, and therefore the prepayments can be estimated at several different sizes at the same node in the tree. This makes it logical to apply another estimation method, and here Monte Carlo simulation is an obvious alternative. Monte Carlo simulation generates various paths for the interest rate, and along the way, prepayments are estimated for each simulated path and at each payment date. For the same reason, Monte Carlo simulation

[74]Sometimes the pool factor have been defined slightly differently, namely as the ratio of outstanding debt to the maximal outstanding debt. The two measures are not too different for relatively newly issued annuities, but will become more and more divergent as time passes. See e.g. Jakobsen & Svenstrup (1999), p. 9.

has been used more and more in the specially designed programs for valuing mortgage bonds in Denmark.

Hence, both maturity and the burn-out effect, measured by the pool factor, are relevant variables that are potentially fruitful to include in the modelling of prepayments.

4.3.3 Loan Size

One thing that may be very important to the decision of whether to prepay a loan or not, is the size of the loan. Since there are transaction costs when prepaying a loan, and subsequently taking on a new loan, and since these costs are in part fixed in nature, obviously sufficiently small loans will not be prepaid, even if the difference between the coupon rate on the existing loan and the refinancing rate is large. So, it is natural to expect that the size of the loan should be positively correlated with the prepayment level. The reason for this is mainly twofold:

- **Transaction Costs**. The existence of fixed transaction costs naturally makes it more profitable to prepay large loans compared to small loans.
- **Clientele effect**. It is natural to expect that mortgagors with large loans are more observant to changing market conditions, and therefore prepay their loans fairly early after the opportunity of gaining from doing so, arises.

The costs of prepaying a loan consists primarily of direct costs to the mortgage bank and direct costs in the form of government taxes. We assume that, if we for a moment disregard other factors, holding interest rate differentials etc. constant, we can write the costs of prepaying a loan as an affine function of the loan size.

$$C(\text{Loan Size}) = \alpha + \beta \cdot \text{Loan Size} \tag{4.9}$$

The gain of prepaying a loan, again for a given interest rate differential, is on the other hand a constant fraction of the loan size. This means that there must be a threshold size, for which we can say that mortgages of a size below this critical value *should not* be prepaid, while mortgages of a size beyond this critical value *may* be prepaid. This is illustrated in Figure 4.3.[75]

Hence, both the transaction cost effect and the clientele effect dictate that large loans should be prepaid before small loans. Another way to express this is that the share of

[75] The figure is obviously drawn for an interest differential that leads to a gain of prepaying – i.e. the bond trades at a rate above par.

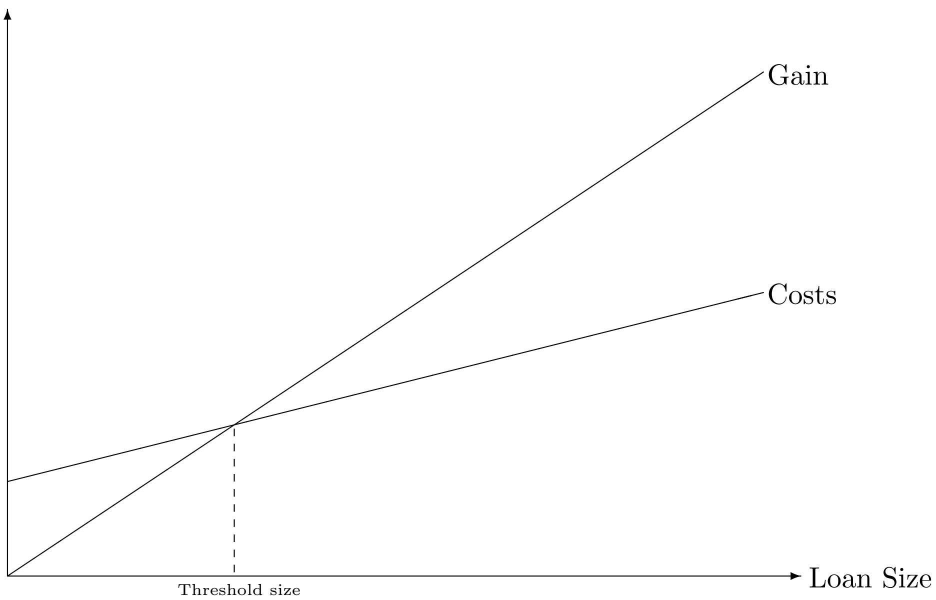

Figure 4.3: Relation between gains / costs of prepaying a mortgage and the loan size.

large loans in a given series should constitute a smaller and smaller fraction of the total loans as time passes. In Figure 4.4, the development in the relative amount of small size loans (less than DKK 500,000), medium size loans (more than DKK 500,000, but less than DKK 3 million) and large loans (more than DKK 3 million) is illustrated for RD 6% 2026.

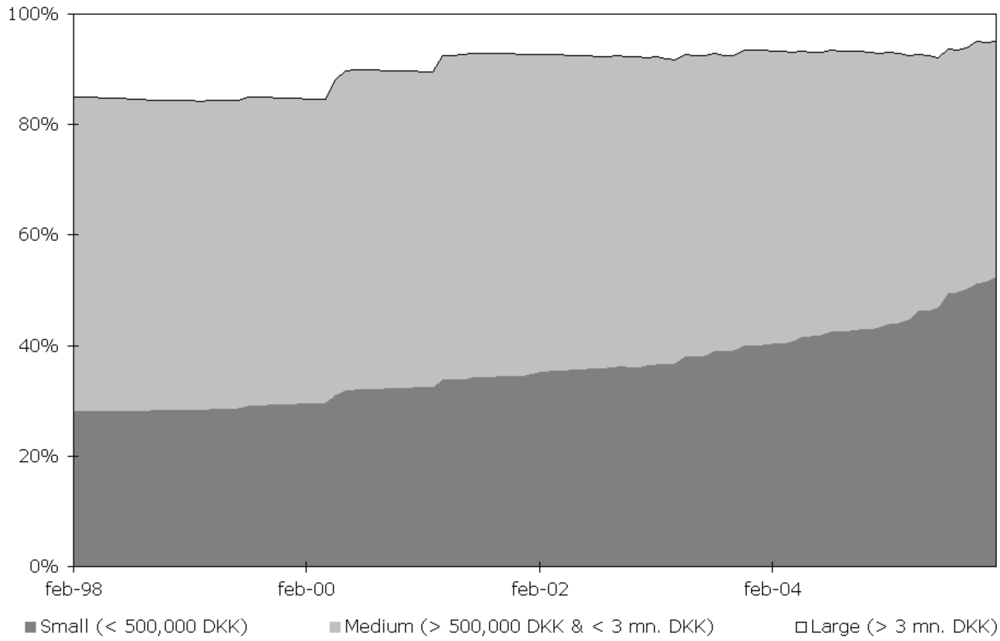

Note: Shares are calculated as shares of outstanding debt.
Source: HSH Nordbank Copenhagen Branch.

Figure 4.4: Loan sizes in RD 6% 2026

This series is associated with a significant prepayment extent since 2002. During this period, it is seen from the figure that the share of small-sized loans is increasing as time passes. This is consistent with the theory of large loans being prepaid before small loans. Note however, that one should in principle take into account the natural ”migration effect” through the loan size groups, since all of the loans, which are annuities, will automatically have a decreasing outstanding debt due to the running payment of instalments. Hence, this figure cannot be taken as proof of this line of theory – that prepayment extent is positively correlated to the loan size, but it is an indication that this line of reasoning may be heading in the right direction, and that it could be fruitful to include some measure of loan sizes in the prepayment model.

Unfortunately, one cannot directly from available market data see exactly which loans are prepaid. One can see the average size of the loans and the amount in the various loan size categories, but this is somewhat approximative. To obtain the exact prepayments in various loan size categories requires a special extract from the database of a mortgage bank. BIS (2004) has obtained access to such an extract from Nykredit, and it confirms that large loans are prepaid significantly faster than small loans.

Even though the loan size seems to be significantly related to prepayments, the loan size is rarely included when setting up a prepayment model. This is due to the fact that the positive relation between prepayments and loan size can to a large extent be captured by an economic gain variable. If, however, the economic gain variable included in the model is an approximative one, such as $\frac{c}{r}$, the inclusion of the loan size as an independent parameter is easily justified. On the other hand, it is worth noting that several models estimate separate prepayment rates for different loan size categories, which also tends to make the inclusion of the loan size as an explicit explanatory variable in the prepayment function superfluous.

In the above sections, we have presented three important drivers of prepayment behavior. Obviously, there may be a list of other factors that may also influence the prepayment level. We save the discussion of additional prepayment drivers and their inclusion into a prepayment model for section 5.4. Before that, we turn towards the modelling of prepayments.

5 Modelling Prepayments

After having discussed variables that are potentially important to prepayment behavior, we now address the question of how to set up a prepayment model. To make things clear, a prepayment model, or equivalently a prepayment function, is a function that gives e.g. the conditional prepayment rate (CPR) – the share of the remaining mortgagors in a series that chooses to repay their loan at a given term, as the output. CPR is then, in the prepayment function, given as a function of various variables, as the ones stated in the previous subsections.

The big issues of making a prepayment model, are what the functional form should look like, which variables to include, and how to estimate the parameters. There is an infinite number of opportunities to choose from here, and different researchers have come up with different specifications, and yet none have been deemed superior. It is fair to say that the functional form, the included variables, and the parameter estimates of the prepayment model is of high importance for the modelling of mortgage bonds. Therefore, the construction of the prepayment model is often a business secret to the participants in the market, due to the importance of this choice. However, it is sometimes public information which variables are included in the prepayment model, but it is usually private information how these models are interacted in the prepayment function; neither is it public information what values the parameters attain. In the following, we address how a prepayment function is modelled in a paper treating the American case and a paper treating the Danish case in sections 5.1 and 5.2, respectively. Afterwards, we turn to the issue of setting up our own prepayment model in section 5.3. Finally, we discuss potential model improvements in section 5.4.

5.1 Goldman Sachs' US Model

An interesting insight into a specific prepayment model is provided in Richard & Roll (1989). The prepayment model they treat is the proprietary prepayment model of Goldman Sachs for the US mortgage security market.[76] This makes the model particularly interesting, since it has been used in the market. The paper states four factors, which are included in the Goldman Sachs prepayment model. These are

– Refinancing incentive

– Age of the mortgage

[76] At least at that point in time – 1989.

– Month of the year (seasonality)

– Burn-out

As mentioned in section 4.3.1, the refinancing incentive can to some degree be accounted for by including the variable $\frac{c}{r}$, such that when the coupon rate is significantly higher than refinancing rate, prepayments would be higher. In other words, the prepayment function is increasing in the fraction $\frac{c}{r}$.

The next factor included in the Richard & Roll (1989) model is the age of the mortgage. They argue that the prepayment rate on average increases over time, even if the interest rate also has increased. This is due to the absence of a delivery option in US mortgage securities. The seasoning or age effect is definitely important in the American case, but less important in the Danish case.

Furthermore, they include the month of the year as a variable as well to incorporate that the turn-over of houses is higher in the summer than in the winter. This variable is rarely included in Danish prepayment models. One argument why the month of the year should not be included in a model for the Danish mortgage bond market, is provided by Figure 4.1, where a seasonal effect does not seem to be prevalent. Possibly, the house turnover effect that admittedly exists, is mitigated by mortgagors' lesser focus on financial issues at this time of year. The latter effect is supported by anecdotal evidence from the Danish mortgage sector.

The last effect they include is the burn-out. This is to account for the effect mentioned in section 4.3.2 that there seems to be a tendency for prepayments to slow down over time, if the series previously has been subject to prepayments.

These are the four variables included in the Goldman Sachs model. The next interesting feature of this prepayment model, is of course how the variables are interacted. This is done by

$$\begin{aligned}\text{CPR} &= (\text{Refinancing Incentive}) \cdot (\text{Seasoning multiplier}) \cdot \\ &\quad (\text{Month multiplier}) \cdot (\text{Burn-out multiplier}) \qquad (5.1)\end{aligned}$$

The multiplicative structure of (5.1) makes the model very sensitive to extreme values. On the other hand, it is an attractive feature that a sufficiently low value of the refinancing incentive variable results in close to zero prepayments, disregarding all other factors.

Through the above model, they are able to explain almost 95% of prepayments over a ten year time span from 1979 to 1988 in the US market. This is an impressive result of a global[77] prepayment model, and they even show that the model can be further improved, as measured by its explanatory power, by separating mortgages into groups of limited coupon rates and/or maturities.

The Goldman Sachs model is an interesting baseline case, both due to the early time of creation of the model and due to the excellent explanatory power that this model seems to exert. It is, however, to be regarded as a baseline case, since we want to focus more specifically on the Danish market. Therefore, we turn our view towards another model that is targeted specifically at the Danish case in the coming section.

5.2 FinE Model

FinE is a Danish developed function library that includes a pricing model for callable Danish mortgage bonds. Hence, it also includes a prepayment model, which is described in Madsen (2005).

Madsen (2005) notes that, in Denmark, the prepayment incentive is solely interest-rate dependent, which is why this is the primary explanatory variable.

The FinE model belongs to the required gains model class, and the gain is measured in relative net present value terms as

$$\text{Gain}^i = \frac{NPV_{old}^i - NPV_{new}^i}{NPV_{old}^i} \tag{5.2}$$

where the net present values are after-tax present values. The superscript i indicates that the prepayment gain is estimated for each loan size group separately. This gain variable defined in (5.2), is obviously an important determinant in the estimation of prepayment rates.

In order to calculate the prepayment behavior in the FinE model, a truncated normal distribution, where the lower tale is cut away at some given level, is applied. The formula for the prepayment rate in the FinE model is then given by

$$\text{CPR}_i = \max\left[\frac{N(G_i;\mu_i,\sigma_i) - N(Level;\mu_i,\sigma_i)}{1 - N(Level;\mu_i,\sigma_i)};0\right] \tag{5.3}$$

which is simply the formula for a truncated normal distribution. The density and distribution functions are depicted in Figure 5.1.[78] From the figures, it is seen that if the

[77] Meaning that the parameters estimated in the prepayment model are the same for all kinds of mortgages, e.g. all coupons, loan sizes, maturities etc.

[78] With more or less arbitrarily chosen values of parameters μ_i, σ_i and $Level$.

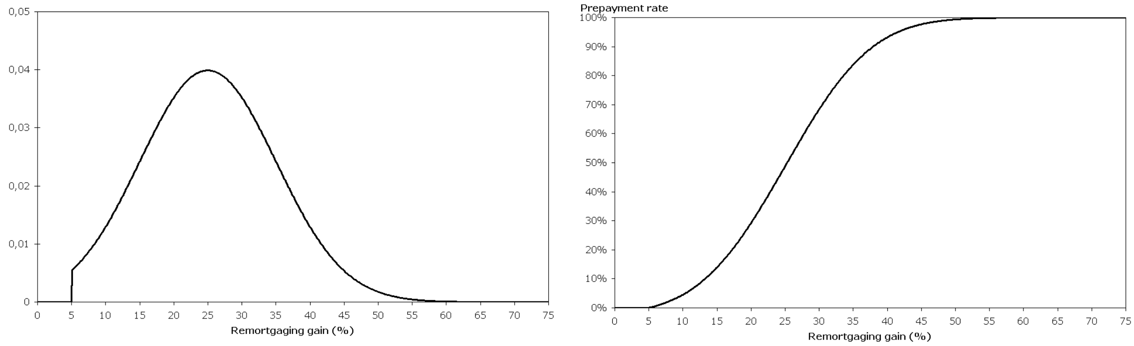

Note: The figures are calculated for $\mu_i = 20$, $\sigma_i = 10$ and Level = 5.

Figure 5.1: Truncated normal distribution

gain from remortgaging (prepaying) a loan is sufficiently small, it is assumed, when applying a truncated normal distribution, that no one will do so. There are good reasons to use a *truncated* normal distribution to model prepayments, since it is not logical to have a prepayment function that predicts, however small, prepayments when the refinancing gain is negative, in the Danish case. This will obviously be the case if an ordinary (non-truncated) normal distribution is used instead.

Actually, (5.3) only defines the shape of the prepayment function – the distribution. The challenge is of course to set up the function G_i that determines the fraction of mortgagors that prepay their loans. Furthermore, values of μ_i, σ_i and *Level* must be determined. The issue of determining an optimal μ_i and σ_i is not dealt with in Madsen (2005), while the estimation of *Level* is simply done by including it in the estimation.

In Madsen (2005), the functional form is not presented due to secrecy, since this is a significant competition parameter in this industry. However, the variables included in the function G_i are described:

$$G_i = f\left(Gain_i, PoolFactor_i, YCSlope, YCChange, \frac{Maturity_{remaining}}{Maturity_{max}}\right) \tag{5.4}$$

Note first that the prepayment function is made path dependent by the inclusion of the pool factor in the prepayment function just like in Richard & Roll (1989). The maturity is included as a fraction of the original maturity of the loan, rather than including it as simply the absolute time to maturity. Madsen (2005) finds the maturity ratio to be positively correlated to the size of prepayments. This is a logical conclusion since the prepayment incentive rises with the maturity ratio.

The new variables, compared to the other models we have seen, are the slope of the yield curve and the change in the yield curve (measured as the change in the refinancing rate from the last term to the present one).

The slope of the yield curve is found to have a positive effect on prepayments. This is explained by what is termed the *FlexLoan* effect. This is due to the fact that a very steep yield curve increases the incentive for mortgagors to shift from long-term traditional mortgage loans to short-term adjustable rate mortgages. When the short-term interest rate is much lower than the long-term interest rate, a lot of mortgagors will be tempted to replace their long-term fixed interest rate loans with short-term adjustable rate mortgages. Hence, this effect is found to be prevalent in the data set under consideration in the paper.[79]

The change in the yield curve must a priori be expected to have a negative impact on prepayments, such that a parallel shift in the yield curve downwards provides a positive impact on prepayments. This effect is confirmed in Madsen (2005).

It is natural to expect the last variable, the pool factor, to have a positive impact on prepayments. This means that as the pool factor falls (as prepayments occur), there should be a dampening effect on prepayments, which means that the pool factor should have a positive contribution to prepayments. This hypothesis is confirmed and furthermore, it is noted that the pool factor contributes to the explanation of the prepayments in a non-linear manner.

The findings in Madsen (2005) are very satisfactory, since the effects from the variables included are just as expected a priori. This, combined with the fact that the model has a very nice explanatory power ($R^2 = 87\%$ in the selected sample), gives us an indication that this way of modelling prepayment behavior, is a reasonable path to follow. The FinE prepayment model is merely described as a "black box" in the sense that the paper does not provide us with any insights as to how the model is actually formulated, but only with a list of the included variables. In the next section, we set up our own prepayment model in order to explain the methodology in dealing with prepayment modelling in practice, both when it comes to the choice of variables and their practical inclusion into the model. This is primarily done to show the considerations that is made when setting up such a model, again emphasizing the practical issues involved in the model set-up.

[79] Even though the positive effect of the yield curve slope on prepayments is termed the FlexLoan effect in Madsen (2005), this effect is of course also relevant for other mortgage loans based on the short end of the yield curve, such as the capped floaters and floating-to-fixed loans. We return to the issue of new products on the Danish mortgage bond market in section 9.

5.3 Our Prepayment Model

As it appears from the preceding sections, there are many things to consider when setting up a prepayment model. The most important issues are

- Which variables should be included?
- What should the functional form be?
- How to estimate the parameters of the model?

In the following, where we go through the construction of a prepayment model leading to the set-up, estimation and evaluation of our own prepayment model, we draw on the findings from the knowledge we have obtained in the previous sections.

We start by addressing the first issue, the variables. The choice of which variables and how many variables to include in the model is, as always, a choice between parsimony of the model and a satisfactory explanatory power. Initially, we select three different variables that we, a priori, think are important drivers of prepayments, drawing on the findings from section 4.3.

Obviously, the first and most important variable would have to cope with the economic gain of prepaying loan. We draw here on the findings in section 4.3.1, and like Richard & Roll (1989), we choose the economic gain variable as an approximative one, namely $\frac{c}{r}$. This approximative variable seemed to be rather good at explaining prepayments, and it is much more simple to include in the model than a more precise NPV variable. We choose not to consider costs of prepayment, knowing that we of course simplify things considerably making this assumption.

One question immediately arises; how should the refinancing rate r be measured? Again, we will have to use an approximation. We choose to use benchmark mortgage rates published by Danmarks Nationalbank, available for maturities of 10, 20 and 30 years. Hence, we choose to use that of the three refinancing rates that is the closest to the remaining maturity as possible, meaning that a loan with 13 years remaining to maturity will use the 10 year benchmark rate as its refinancing rate, while a loan with 27 years remaining to maturity will use the 30 year benchmark rate as its refinancing rate. Another approach could be, like in Madsen (2005), to make skew intervals, based on an expectation of upward biased refinances.[80] However, we choose to use the simple

[80]Since to lengthen the time to maturity is an easy way to decrease the first year payment.

non-skew measure, such that the refinancing rate is

$$r(\text{Maturity}) = r_t^{\text{benchmark}}, \qquad t = \text{Round}(\text{Maturity}/10; 1) \cdot 10 \tag{5.5}$$

Drawing on the arguments from section 4.3.2, we also include the maturity as an explanatory variable, defining the maturity variable simply as the remaining time to maturity. We choose not to include the pool factor in the model.

It is a normal way to proceed, to estimate prepayment rates for different loan size groups separately, based on the fact that different loan sizes prepay at different speeds. We choose not to follow this path, but instead include data for the average loan size in the estimation. In this way, we should be able to separate the loan size effect (if any) on prepayments more easily. However, the measure may also prove to be too simple to have any explanatory power.

To sum up, these are the three variables that we choose to include in our first model:

- The **refinancing incentive**, measured by $\frac{c}{r}$.
- The **time to maturity**, measured in years.
- The **average loan size**.

Next, we have to decide on a functional form. As the functional form in the estimation, we use a probit formulation following Jakobsen (1992). The use of a probit funtion is justified by the intuitive attractiveness of an underlying normal distribution dictating the propensity to prepay. This means that the prepayment function, i.e. the function giving CPR as output is given by

$$CPR = N\left(a_0 + \sum_k a_k \cdot x_k; \mu, \sigma\right) \tag{5.6}$$

where the x_k's denote the three input variables. With the chosen variables, the function can be written as

$$CPR = N\left(a_0 + a_1 \cdot \frac{c}{r} + a_2 \cdot Maturity + a_3 \cdot LoanSize; \mu, \sigma\right) \tag{5.7}$$

We transform this equation in order to obtain a probit model based on the standardized normal distribution. This is simply done by subtracting the mean and dividing by the standard deviation. Hence, we write

$$CPR = \Phi\left(\beta_0 + \beta_1 \cdot \frac{c}{r} + \beta_2 \cdot Maturity + \beta_3 \cdot LoanSize\right) \tag{5.8}$$

where $\beta_k = \frac{a_k}{\sigma}$ for $k > 1$ and $\beta_0 = \frac{a_0 - \mu}{\sigma}$.

To estimate this non-linear function, we choose to apply the maximum likelihood estimation method.[81] We define a variable y_{imt}, which is a binary variable, indicating the choice of mortgagor i in pool (bond series) m at time t of whether to prepay his loan ($y_{imt} = 1$) or not ($y_{imt} = 0$). We therefore see the function for CPR as

$$P[y_{imt} = 1] = \Phi\left(\beta_0 + \sum_k \beta_k \cdot x_{kmt}\right) \tag{5.9}$$

The one-dimensional stochastic variable y_{imt} follows a binomial distribution with the time and bond varying event probability parameter $\theta_{mt} = P[y_{imt} = 1]$, which is itself a stochastic variable that is normally distributed and given by (5.9). Since y_{imt} follows a binomial distribution, we can write the density function for this variable as

$$f_{y_{imt}}(y_{imt}, \theta_{mt}) = \theta_{mt}{}^{y_{imt}} \cdot (1 - \theta_{mt})^{(1 - y_{imt})} \tag{5.10}$$

If we assume that the observations are independent, we can write the joint density as the product of the individual densities.[82] Here, we skip the subscript i, reducing the three-dimensional stochastic variable $\mathbf{y}$ to be a two-dimensional variable, since we only have aggregate data – not on level of individuals.

$$f_{\mathbf{y}}(\mathbf{y}, \theta) = \prod_{m=1}^{M} \prod_{t=1}^{T} \theta_{mt}{}^{y_{mt}} \cdot (1 - \theta_{mt})^{(1 - y_{mt})} \tag{5.11}$$

The concept of maximum likelihood is that the joint density function is set up, and the probability of observing the actually observed data set is then maximized. In that way, the estimated parameters are the ones that give the specified function the highest probability of observing what was actually observed. So, the likelihood of observing the data set, given the model at hand, is maximized. We can easily find the likelihood function from

[81] For background references on maximum likelihood, see e.g. Gabrielsen, Kousgaard & Milhøj (1999) or Verbeek (2004), chapter 6 & 7. We only go briefly through the concepts here.

[82] This assumption is problematic, as the CPR's cannot be expected to be independent. However, it is a common approach used in the literature to maintain this assumption.

the joint density function (5.11)

$$\mathcal{L}(\theta|\text{CPR}) = \prod_{m=1}^{M}\prod_{t=1}^{T}\theta_{mt}{}^{\text{CPR}_{mt}} \cdot (1-\theta_{mt})^{(1-\text{CPR}_{mt})} \tag{5.12}$$

where the matrix CPR of observed quarterly prepayment rates, replaces the stochastic variable $\mathbf{y}$. Inserting the expression for θ_{mt} gives us the following expression, where θ_{mt} is replaced by the expression in (5.9)

$$\mathcal{L}(\theta|\text{CPR}) = \prod_{m=1}^{M}\prod_{t=1}^{T}\Phi\left(\beta_0+\sum_k \beta_k\cdot x_k\right)^{\text{CPR}_{mt}} \cdot \left(1-\Phi\left(\beta_0+\sum_k \beta_k\cdot x_k\right)\right)^{(1-\text{CPR}_{mt})} \tag{5.13}$$

This is the function that should be maximized with respect to the β's. This is done by differentiating the likelihood function with respect to the parameters. Normally, it is much easier to maximize the logarithm of the likelihood function, since products are then transformed into sums. This is a valid transformation, since the logarithm is a positive monotonous transformation. Making this transformation (5.13) provides us with the following log-likelihood function

$$\log\mathcal{L}(\theta|\text{CPR}) = \sum_{m=1}^{M}\sum_{t=1}^{T}\Bigg[\text{CPR}_{mt}\cdot\log\Phi\left(\beta_0+\sum_k \beta_k\cdot x_k\right) + (1-\text{CPR}_{mt})\cdot\log\left(1-\Phi\left(\beta_0+\sum_k \beta_k\cdot x_k\right)\right)\Bigg] \tag{5.14}$$

Since there is a vector of parameters, the first derivative of (5.14), which is known as the *score vector*, is also a vector. The score vector should be calculated and then set equal to the zero vector. The function (5.14) is, however, very hard to differentiate, and therefore we use a computer program to maximize it numerically instead of analytically. We choose to apply the SAS® package.[83]

Before we do this, we have to ensure that when maximizing the log-likelihood function, we actually find a *global* maximum. Fortunately, Amemiya (1981) argues that this log-likelihood function is globally concave, given that all prepayment rates lie between zero and one. Therefore, we do not need to feed the software with initial parameter estimates.

The sample, which is used for the estimation, is based on a list of 144 different mortgage bonds from various mortgage banks with coupon rates of 5, 6, 7 and 8% in the period from the January term 2001 to the January term 2006. This gives us a sample of 2838 quarterly

[83]The PROBIT procedure is employed; the necessary programming lines are shown in Appendix B.2.

Parameter	Variable	Estimate	Std. error	p
β_0	Constant	-2.8486	0.2617	0.0000
β_1	$\frac{c}{r}$	0.9182	0.1311	0.0000
β_2	Maturity	0.0264	0.0056	0.0000
β_3	Loan Size	-0.0000	0.0000	0.2199

Source: Own calculations conducted in SAS® on data obtained from HSH Nordbank Copenhagen Branch.

Table 5.1: Parameter estimates in prepayment model #1

observations in total. The estimation of the prepayment model yields the parameter estimates shown in Table 5.1.

In order to test the significance of the included variables, we apply a Wald test. The Wald test principle is to estimate the parameters only under the alternative hypothesis, and to check whether the ML-estimate is significantly different from the null hypothesis value using the estimated covariance matrix. The Wald test size for a hypothesis of $\mathbf{A}\beta = \mathbf{c}$ is given by[84]

$$W = (\hat{\beta} - \mathbf{c})^\intercal \left(\mathbf{A}(-\mathbf{Q}(\hat{\beta}, n))^{-1}\mathbf{A}^\intercal\right)^{-1} (\hat{\beta} - \mathbf{c}) \sim \chi^2(r) \tag{5.15}$$

where $\mathbf{A}$ is a $r \times p$ matrix, where r is the number of restrictions and p is the number of parameters in the model. The test size is a quadratic form, where the mid section of (5.15) is an estimate of the covariance matrix. In the simple case of only one restriction, e.g. testing for significance of one of the variables, the matrix $\mathbf{A}$ turns into a p-dimensional vector, and the quadratic form test size becomes one-dimensional

$$W_k = \frac{(\hat{\beta}_k - 0)^2}{\widehat{\text{var}}(\hat{\beta}_k)} = \frac{\hat{\beta}_k{}^2}{\widehat{\text{var}}(\hat{\beta}_k)} \sim \chi^2(1) \tag{5.16}$$

The significance levels of the included variables calculated from the chi-square distribution with one degree of freedom, are shown in the last column of Table 5.1.

We go through the parameters one by one. The first variable is the economic gain variable, $\frac{c}{r}$. This variable has a positive impact on prepayments. This is as expected, since the higher the coupon rate is relative to the refinancing rate, the higher the fraction of mortgagors must be expected to prepay their loans. The effect of this variable is

[84]We refer to Johnston & DiNardo (1997) for a thorough description of the three different test principles under the maximum likelihood estimation method.

furthermore highly significant, completely as expected.

The next variable is the time to maturity. This effect also has a positive effect on prepayments, such that the longer the time to maturity, the higher the likelihood of prepayments occurring. This is also completely as expected, cf. the discussion in section 4.3.2. This variable is also highly significant.

The last variable is the average loan size. A priori, we expected this variable to have a positive impact on prepayments, such that the mortgages in a series with higher loan sizes would prepay faster than the mortgages in a series with lower loan sizes. This effect is not found in this specification of the prepayment model with this particular sample. The effect is surprisingly found to be negative, but it is not significantly different from zero, cf. Table 5.1.

Since the average loan size does not contribute to explain prepayments in the present set-up, we choose to remove this variable from the model. On the other hand we try to include two new variables, inspired by Madsen (2005).

The first new variable is the slope of the yield curve. The reasoning behind including this variable in the estimation is that the higher the slope of the yield curve, the higher is the incentive to convert a traditional long-term loan into a short-term adjustable rate mortgage. Therefore, we expect that a steep yield curve results in a higher level of prepayments, all other things equal.

The second new variable that we include is the change in the refinancing rate from the last period to the actual. Here, we simply calculate the absolute change (percentage points) in the benchmark refinancing interest rate from the previous exercise date to the actual exercise date. We expect that a change upward in the refinancing rate causes a decrease in prepayments, and a change downward in the refinancing rate causes an increase in prepayments. Of course, this effect is primarily captured by the economic gain variable, but there may be something left to explain for the change in the interest rate variable. This may be due to some sort of a momentum effect; such that rising interest rates inherently cause prepayments to decrease, even though the gains from prepaying may still be significant. Obviously, the interest rate change variable would have little relevance in a rational prepayment set-up, so the effect from this variable is solely driven by the change in the sentiment; the dampened enthusiasm concerning prepayments or vice versa.

Now, we estimate the new model, which consists of the following four variables:

– **The refinancing incentive** $\frac{c}{r}$

Parameter	Variable	Estimate	Std. error	p
β_0	Constant	-3.4337	0.2548	0.0000
β_1	$\frac{c}{r}$	1.0488	0.1172	0.0000
β_2	Maturity	0.0079	0.0068	0.2407
β_3	YCSlope	0.3011	0.0718	0.0000
β_4	Δr	-0.3630	0.1195	0.0024

Source: Own calculations conducted in SAS® on data obtained from HSH Nordbank Copenhagen Branch.

Table 5.2: Parameter estimates in prepayment model #2

– **Time to maturity**

– **Slope of the yield curve**

– **Change in refinancing interest rate**

The results of the estimation of the probit function based on these four variables are shown in Table 5.2.

In this new model, the refinancing incentive variable $\frac{c}{r}$ is still positive and highly significant. The maturity variable is also still positive, as expected. However, with the exchange in the list of variables, it now becomes insignificant, since the p-value of a test that the effect of the maturity variable is actually zero, is around 24% as indicated in Table 5.2. This change indicates correlation between the explanatory variables in the models, but this is not to be regarded as a major problem, since the primary issue of this exercise is more the prepayment predictions rather than the partial effects.

The first of the two new variables included, the slope of the yield curve, has a positive impact on prepayments, exactly as we expected it to have.[85] So, the steeper the yield curve, the higher the amount of prepayments, since the incentive to replace a long-term fixed interest rate loan with a short-term adjustable rate mortgage, becomes higher. The impact of the yield curve slope variable is furthermore highly significant.

The second new variable that was included was the change in the refinancing interest rate from the last exercise date to the actual exercise date. The effect of this variable is negative, which is also completely as expected, such that a rise in the refinancing interest rate from the previous exercise date will, all other things equal, cause a lower amount of prepayments in the actual period. Not only is the sign of the effect of this variable as we

[85]This variables is probably most relevant for a bullish steepener, that is, where the short end of the yield curve drops more than the long end of the yield curve.

expected it to be, it is also significant.

Now we have an idea of which variables make a significant contribution to the explanation of prepayments and in which direction. In the second model, apparently all the included variables, except for the maturity, are significant drivers of prepayment behavior in Denmark in the observed period. One could argue that the maturity should be left out of the model as we did it with the average loan size going from the first to the second model. However, we refrain from doing so, asserting that the found positive affect (however, statistically insignificant) is to be regarded as a relevant driver of prepayments. We admit that the other three variables are statistically much more significant, and they may also be more important factors than the maturity when explaining the prepayment behavior.

An interesting implication of the findings in this section, is that the important drivers of prepayments seem to be almost exclusively interest rate driven. Even though we argue that the maturity should stay in the prepayment model in spite of its apparently insignificant contribution to the explanation of prepayments, we note that the three significant variables in the second prepayment model are all very much interest rate based: The fraction between the coupon and the refinancing interest rate, the slope of the yield curve, and the change in the refinancing rate from the previous exercise date to the actual. When modelling prepayments, we are fully aware that what we do, is in fact to model human behavior. A fact that immediately poses difficulties when trying to create a reliable global model. However, it appears as if the mortgagors in Denmark are (perhaps surprisingly) rational in their prepayment decisions, since, in our treatment, the significant variables are, as we mentioned, based on properties of actual interest rates.[86] On the other hand, the more indirect variables such as the average loan size and the maturity do not seem to play a major role in explaining prepayments.

The next issue is of course to evaluate the explanatory power of the estimated model. This is a question of comparing the predicted values of the prepayment rates with the actual values of the prepayment rates. In Figure 5.2, a plot of observed prepayment rates against predicted prepayment rates is shown.

A natural way to compare estimated and observed values of the prepayment rate is to apply ordinary linear regression methods. Therefore, we apply linear regression to the data set shown in Figure 5.2. Doing this, the optimal situation would be that all the observations are found on a 45° line that goes through the origin, since this would cause

[86]However, rationality is somewhat limited, since the momentum effect should obviously not be prevalent in a rational prepayment set-up.

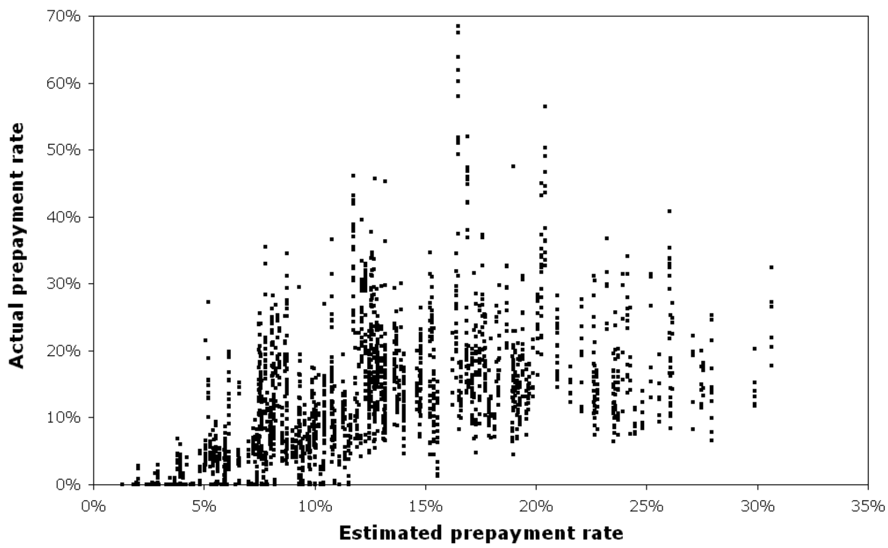

Source: Own calculations conducted in SAS® on data obtained from HSH Nordbank Copenhagen Branch.

Figure 5.2: Plot of observed and estimated prepayment rates

Slope	Std. error	R^2
0.9701	0.0114	71.74%

Source: Own calculations conducted in SAS® on data obtained from HSH Nordbank Copenhagen Branch.

Table 5.3: Parameter estimates in the origo regression based on the data in Figure 5.2

estimated and actual prepayment rates to correspond one to one on average. Furthermore, going through the origin would ensure that there are no systematically added bias. A preliminary regression shows that the vertical intercept is indeed very close to zero.[87] Hence, we follow Madsen (2005) and apply an origo regression to the data shown in Figure 5.2. We would like for this regression line to have a slope as close as possible to unity, and obviously as high an R^2 as possible. The slope and its standard error along with the R^2 is shown in Table 5.3.

It is seen that the slope parameter is fairly close to one; actually so close to one that the average systematic errors are acceptable. Hence, the model does not seem to overestimate or underestimate the prepayment rates dramatically.

With the usual precautions for using R^2 as a measure of explanatory power of the model, we apply it to facilitate comparisons to e.g. Madsen (2005). It is not straightforward to construct an R^2 measure in this case where we have a regression without an

[87] The intercept is estimated to 0.07%, and a test for a hypothesis that the intercept is actually zero cannot be rejected, since the p-value of this t-test is as high as 83%.

intercept. Hawkins (1980) and Kvålseth (1985) provide good insights on the issue of regression without an intercept term, and the latter in particular treats the issue of how to correctly calculate an R^2 in regressions without an intercept term. Kvålseth (1985) argues that in case one leaves out the intercept term in a regression, one should replace the usual R^2 defined by $R^2_{ordinary} = 1 - \frac{\sum(y-\hat{y})^2}{\sum(y-\bar{y})^2}$ with another measure of R^2 defined by $R^2_{origo} = 1 - \frac{\sum(y-\hat{y})^2}{\sum y^2}$.[88] The R^2 calculated by the redefined formula is 0.7174 as seen from Table 5.3. The choice of this redefined R^2 facilitates a direct comparison to the FinE model, where an R^2 of 0.8688 was obtained. However, the redefined R^2 is hardly interpretable as a percentage part of explanatory power out of the total variation in the data. Glancing a moment at Figure 5.2, it also seems surprising that the model at hand should explain 71% of the total variation in the prepayment rates. There is quite a lot of variation left to explain, which would require using more variables, and perhaps consideration of other functional specifications. However, if we take our model's simplicity into account, we must say that an R^2 of 0.7174 is very satisfactory. It is comforting to know that such a simple model can provide with such a good result, and it provides us with the incentive to look further into how the model could be improved. In the next section, we discuss how one could potentially improve the model by introducing new variables.

5.4 Model Improvements

Even though we concluded that the prepayment model developed in the previous section had a satisfactory predictive power, we still find it important to outline how attempts to improve the model further could be done. We look into a list of factors not included in the model so far, which we believe to be potentially relevant variables in the modelling of prepayments.

Looking at Figure 5.2 once again, there appears to be some kind of clustering in this figure. This is to be understood such that it seems that for each estimated prepayment rate, there are about a handfull of actually observed prepayment rates attached to it. This is due to the fact that we have, in the sample, included observations from different mortgage banks with otherwise identical properties in order to make the model more robust, i.e. by including more observations. The clustering pattern is particularly easy to observe where the density of observations is rather low, e.g. for high estimated prepayment

[88] Note that the usual R^2 measure and the redefined one coincide in the special case of a sample mean of zero.

rates.

It would be interesting to see if the prepayment rates are actually significantly different for different mortgage banks on average. This could actually be implemented relatively easily by introducing dummy variables for each (but one) mortgage bank. We have tried to extend the model from the last section with mortgage bank dummies. However, the effect from these variables is very small, and clearly insignificant. Hence, nothing in our data set combined with the model that we have set up suggests any systematic differences of prepayment rates across mortgage banks. The opposite conclusion would also have suggested that there should be systematic price differences between mortgage bonds from different mortgage banks, which we do not observe.

Another improvement attempt could be to try to correct for the *media effect*. The term covers the fact that the prepayment extent may be severely affected by the media effort exerted by the mortgage banks. The last few years, the mortgage banks have been very active in their efforts to make the mortgagors convert their loans whenever profitable.[89] To correct for this factor is very difficult, since one would need an indicative variable of the media effort exerted by the mortgage banks. Even though indexes of e.g. television commercials are publicly available, much of the information of media campaign expenditures is private information to the mortgage banks. It would be very hard to obtain these data, but it would obviously be interesting to include such a variable in the model, provided it could be obtained. However, a vast amount of correlation must be expected between an indicative media effort variable and various e.g. gain variables, since the mortgage banks must be expected to advertise more heavily in periods where prepayments are profitable, i.e. falling interest rates. Therefore, it would be very hard to separate the effect of media campaigns from other variables.

Another thing that could be extremely interesting and very relevant, is to investigate, not just the size of prepayments at different times, but also data on micro level, such that one could see the transition that happens in connection with prepayments. In other words, data indicating not only what mortgagors converted from, but also what they converted into, could be very helpful to help explain prepayment behavior. Unfortunately, such data is not publicly available.

Instead of using such micro data, it might prove useful to include the debtor distributions instead. The information entailed in these data is, among other things, the loan sizes divided into five different size groups. These could be relevant to exploit in order to make a better inclusion of the loan size as a driver than the average loan size, which we

[89]This obviously has to do with the fact that the mortgage banks earn fees every time a loan is prepaid.

unsuccessfully tried to include. Furthermore, the debtor distribution also includes data of the type of debtor, i.e. private or corporate. An idea could be to include the fraction of private debtors as an explanatory variable. This builds on the previously noted observation that corporate mortgagors can be expected to manage their debt more actively and rationally.

The next issue is how to include emersion of new products in the model. The mortgage banks have been very eager to disseminate the knowledge of the latest inventions on the Danish mortgage market, the capped floating rate loans, which we will return to in section 9. One could say that the launch of a new product on the Danish mortgage market could in itself contribute to higher prepayments, since people would convert their loans into a new loan type, even though the gain in the traditional sense, may be very limited, or even negative, if the new loan type was sufficiently attractive. On the other hand, if the loan palette keeps on enlarging, it may also be that people may be reluctant to prepay an existing loan, expecting new attractive loan types to appear in the near future. Thus, the effect of new loan types on prepayment is ambiguous. However, we expect the first of the two effects to dominate.

Furthermore, it could be interesting to see if the expectations to the development in the interest rates could be a driver of prepayments. This is definitely problematic, since in principle, the Hull-White term structure model dictates the expected development in the interest rates. Obviously, when including the expectations to future interest rates, it should not be these interest rates that are taken into account. Rather, it should be the expectations held by the mortgagors. These expectations will most likely be formed by the views expressed by financial advisors, such as commercial banks, central banks, mortgage banks, pension funds etc. An interesting approach could be to include a consensus estimate of future interest rates and see whether this could provide a significant contribution to the explanation of prepayments.

While the previous extensions have received a fair amount of attention in the literature and by industry quantitative research units, the next subject has been somewhat neglected. We now spend some effort on a particular extension to the existing model, namely to exploit preliminary data for prepayments. The idea is to investigate whether the preliminary data can be used as an indicator for the final prepayment extent at the exercise date, thereby enhancing the prepayment model. The value of a mortgage bond is significantly influenced by the prepayment extent at the forthcoming term, and we therefore devote some effort to investigate this issue in the following.

As already noted in section 4.1, mortgagors who wish to redeem their loans have to inform the mortgage bank at least two months in advance. However, the mortgagors inform their mortgage bank that they wish to exercise their prepayment option, also between exercise dates, i.e. before the deadline two months before the given payment date. Hence, the mortgage banks can continuously follow the development in the so-called *scheduled prepayments*. In other words, the scheduled prepayments are the loans that, at a given point in time, have been scheduled prepaid, such that the mortgagors have already (before they actually have to) told their mortgage bank that they want to exercise their prepayment option. It may be perfectly logic for a mortgagor to do so, since the mortgagors can engage in an arrangement with the mortgage bank between two exercise dates, where they make a deal to redeem the existing loan at the next payment date, and at the same time take on a new loan at the **present** interest rate. Since there is no guarantee that the interest rate, i.e. the issue price of the new loan, will be the same at the next exercise date as it is initially, a sufficiently risk-averse mortgagor, or a mortgagor holding the expectation of adverse movements (from his point of view) in the interest rate, may want to make the loan conversion immediately, i.e. before the exercise date. Of course, the mortgagor incurs a cost for this service, but he may still find this to be a good idea.

Hence, the mortgage banks obtain knowledge of scheduled prepayments during the period between two exercise dates. The scheduled prepayments should be expected to be a good signal to the market of how large the extent of total prepayments will be at the next payment date. Therefore, the scheduled prepayments can potentially be used to improve the prepayment prediction for the next term. The information is made public through Copenhagen Stock Exchange, and it is followed closely by market participants. This information is calculated by all mortgage banks in the Danish market every Friday at noon, and is published the following Tuesday at noon.[90]

We now investigate whether it should be possible to observe some pattern in the scheduled prepayments during the time between two exercise dates. It turns out that there is indeed a very solid pattern that is observed in practically all series. The scheduled prepayments for RD 6% 2035 for the last three terms in 2005 are shown in Figure 5.3.

There seems to be a particular pattern in the graphs. All the graphs seem to show a behavior that potentially could be explained by an exponential or a power function. Before we proceed to check whether this is actually the case, note the peculiar kinks, many of the graphs have in the last period. This is due to the irregularity of the length

[90]Madsen (2005), p. 2.

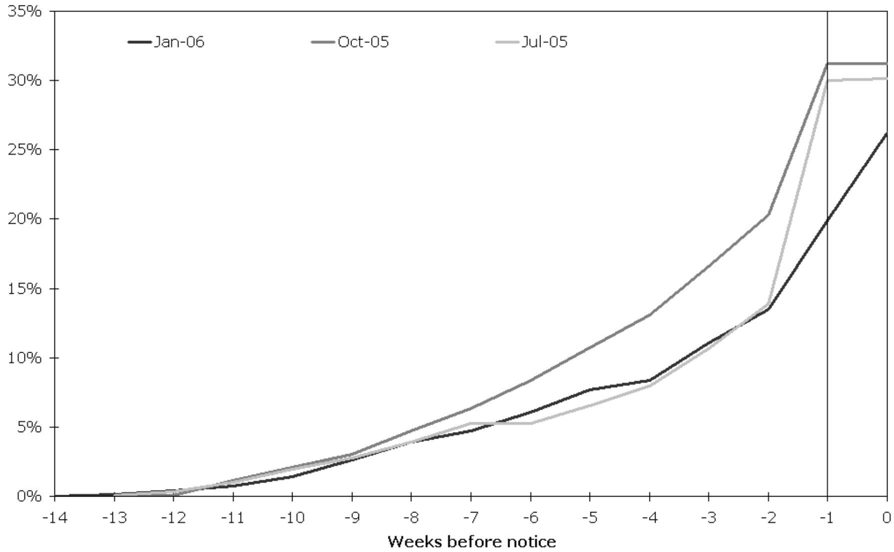

Source: HSH Nordbank Copenhagen Branch

Figure 5.3: Scheduled prepayments for RD 6% 2035

of the last "week".[91] Therefore, if one wants to try to predict the prepayments at the present term, one will have to find out exactly how long the last period is at the present term, in order to make a reliable forecast. To simplify things, we restrict ourselves to estimate the prepayments at the second-to-last observation (time "–1" in Figure 5.3). We also restrict ourselves to looking at this particular bond shown in Figure 5.3, RD 6% 2035, which is chosen more or less arbitrarily among mortgage bonds with a recent significant prepayment extent.

First, we try to see if an exponential regression provides a good fit, since the shape of the curves could, at a glance, seem to have an exponential shape. It quickly turns out that this is definitely not the case. The regression curve for an exponential function by far overestimates the prepayments. This is not so mysterious, since, if we draw the logarithm of the scheduled prepayments as a function of time, the structure is nowhere close to a straight line, which it should be if the scheduled prepayments were to develop according to an exponential function. This is shown in the left part of Figure 5.4. The curve is obviously increasing, but with a decreasing slope, indicating that the exponential functional form is, in a sense, too increasing to capture the behavior of the scheduled prepayments as a function of time. In the right part of Figure 5.4, the logarithm of scheduled prepayments is shown as a function of the logarithm of (a modification of)

[91] As noted, the last exercise of the prepayment option is exactly two months before the payment date. Data deliverances happen on a weekly basis, and, as noted, the data is always updated Friday at noon. Therefore, the last "week" may be of variable length – from half a day to $6\frac{1}{2}$ day.

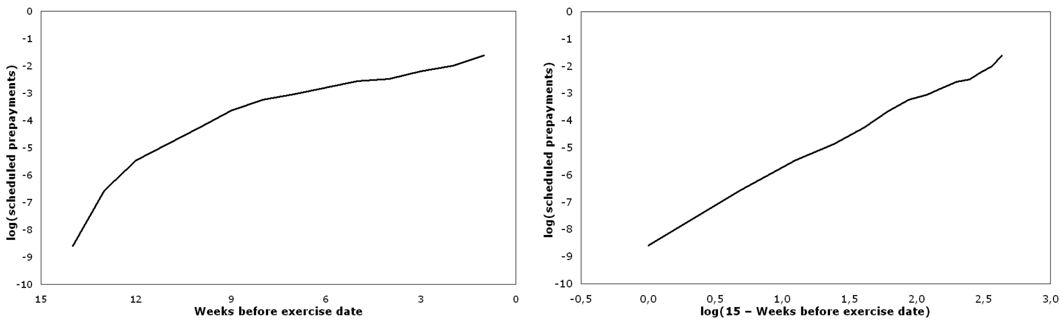

Figure 5.4: Scheduled prepayments for RD 6% 2035, January 2006 term

time. This correspondence looks very linear, leading one to suggest that the connection between time and scheduled prepayments may be some sort of a power function, i.e. a function of the following functional form:

$$\text{Scheduled Prepayments} = b \cdot [15 - \text{Weeks before exercise date}]^a \tag{5.17}$$

So, by estimating (5.17) on basis of the scheduled prepayments announced until a certain point in the period between two exercise dates, one should be able to give a fair estimate of the scheduled prepayments at the exercise date or any other time before this – in our case at the second-to-last week. Let us now investigate further how good the explanatory power of such a simple model is. We do this by once again looking at the scheduled prepayments for the January term 2006, and estimating (5.17) recursively, each time including yet another observation. Expectedly, the estimate should become better as time elapses.

From Table 5.4, it is seen that the model based on a regression of equation (5.17) does a very poor job in the beginning of the period. This is not too bad since this very naive approach should be expected to require some more observations before it is capable of making a decent estimate. Already from the inclusion of five observations, i.e. from ten weeks before the exercise date and onwards, the model actually makes a pretty good estimate up until the exercise date, with a maximum error of 5.25 percentage points, and it is more or less systematically converging towards the actual number, such that the estimate gets better and better as more and more observations are included. This is a fairly nice result of such a simple model. There is, however, a serious problem with the estimation results. They seem to be rather systematically biased. The estimate of the prepayments falls almost through the entire period. This may be sample specific, but could possibly be the result of a model that is wrongly specified. It is not terribly surprising that a model as simple as this one does not possess the whole truth, but it still

Date	Weeks before exercise	Observations	Scheduled Prepayments	$\hat{a}$	$\hat{b}$	Estimate	Error in percentage points
Aug 2, 2005	14	1	0.02%	N/A	N/A	N/A	N/A
Aug 9, 2005	13	2	0.14%	0.000184	2.95	44.55%	24.63
Aug 16, 2005	12	3	0.43%	0.000187	2.87	36.32%	16.40
Aug 23, 2005	11	4	0.77%	0.000196	2.73	26.18%	6.27
Aug 30, 2005	10	5	1.42%	0.000200	2.68	23.88%	3.96
Sep 6, 2005	9	6	2.61%	0.000198	2.70	24.85%	4.93
Sep 13, 2005	8	7	3.91%	0.000197	2.71	25.17%	5.25
Sep 20, 2005	7	8	4.74%	0.000202	2.68	23.63%	3.71
Sep 27, 2005	6	9	6.08%	0.000208	2.64	22.26%	2.35
Oct 4, 2005	5	10	7.67%	0.000213	2.61	21.18%	1.26
Oct 11, 2005	4	11	8.36%	0.000222	2.57	19.69%	-0.23
Oct 18, 2005	3	12	11.04%	0.000228	2.55	18.95%	-0.97
Oct 25, 2005	2	13	13.47%	0.000232	2.53	18.43%	-1.49

Table 5.4: Power regression on scheduled prepayments for RD 6% 2035, January 2006 term

provides the motivation to go further down this path.

In the end, a simple regression model of some sort of a functional form with a very limited number of parameters could not be expected to be able to make a very good estimate of prepayments for the next term based on the scheduled prepayments at a given point in time. The results we have obtained are surprisingly good, but the prepayment extent is ultimately a result of human behavior combined with a list of outside factors. Therefore, it is natural to include many other factors in the estimation of prepayments, also for the coming term, than simply just time, most importantly of course the refinancing incentive.

This naturally calls for the combination of a normal prepayment model, as the one that we have just developed in section 5.3 and the scheduled prepayments. This is actually also the way that some of the models that actually do use the scheduled prepayments, include them. Normally, the industry prepayment models put entirely focus on the traditional prepayment modelling method in the beginning of a period, but as the exercise date approaches, more and more weight is put on the scheduled prepayments. It is fair to say that the inclusion of scheduled prepayments may provide a refinement to an existing prepayment model, and in many cases this could be a significant contribution, as the example we have provided in this section shows. However, to imagine that the scheduled prepayments could replace all other explanatory factors in a prepayment model is obviously naive. Nevertheless, it is interesting to investigate further how much and in what way the traditional prepayment models can be refined with the inclusion of scheduled prepayments.

One path that has been tried taken is the method of neural networks. Basically, neural networks is a way of trying to make a computer see connections in different situations. A keyword for neural networks is learning. You try to teach the neural network to see connections in the situations you wish to analyze. The term *neural networks* comes from the fact that the way you try to teach the network to see patterns, is the same way as the human brain does it.[92]

The neural network learns by experience, which means that it goes through historical data and applies an extensive list of algorithms to the data at hand. By analyzing this data, it is trained to be capable of seeing systematics in the future data. This can be used for prepayments. The network should be trained for making an estimate of prepayments on basis of the pattern of scheduled prepayments at a given point in time. The network then does so by applying the knowledge obtained from the pattern of historical prepayments. Obviously, the network should have access to a wide list of variables that may be explanatory in this connection. There have been few attempts to apply this in the present context, but the attempts made have been very prosperous.[93]

This concludes the treatment of prepayments in this thesis. In the next section we briefly discuss the principle of combining the term structure and prepayment models, before we turn towards investment issues in sections 7 and 8.

[92]Neural networks have overwhelming perspectives; neural networks are often regarded to be a premature step towards the creation of sophisticated artificial intelligence.

[93]To our knowledge, in Denmark, only Nordea has tried to use neural networks for estimating prepayments. The network we have been presented for, was able to explain 84% of the prepayments.

6 Combining Term Structure and Prepayments

Now, we have gone through the issues of modelling the term structure and how to apply this model as well as the set-up of a prepayment model. As illustrated in Figure 1.3 on p. 6 in the introduction, these are the two main necessities needed when setting up a pricing model for callable mortgage bonds. Even though we do not set up a full pricing model in this section, as the technical set-up of a full model is out of the scope of this thesis, we outline the principle of it here.

As shown in section 2, the value of a mortgage bond must be equal to the discounted expected cash flow of the mortgage bond under the martingale measure. This means that the value of the mortgage bond, with a principal of unity for convenience, is given by

$$P(0) = \sum_{n=1}^{N} PV(CF_n) \tag{6.1}$$

Here, the present values of the cash flow at time t_n – i.e. CF_n can be calculated as

$$PV(CF_n) = E^Q\left[e^{-\int_0^{t_n} r_x dx} \cdot CF_n\right] \tag{6.2}$$

Hence, in order to calculate the value of the mortgage bond in its fullest, we merely need an expression for the cash flow of the bond. The cash flow of the mortgage bond consists of prepayments, interest payments and ordinary redemptions. Hence, we can write the cash flow of the mortgage bond at time t_n as

$$CF_n = \underbrace{\left(\prod_{i=1}^{n-1}(1 - CPR_i)\right)}_{\text{Nominal left at time } t_n} \cdot \underbrace{\left(\underbrace{CPR_n}_{\text{Prepayments}} + \underbrace{(1 - CPR_n) \cdot o_n}_{\text{Ordinary redemptions}} + \underbrace{\frac{c}{frq}}_{\text{Interest payment}}\right)}_{\text{Cash flow of nominal 1 at time } t_n} \tag{6.3}$$

Here, o_n denotes the ordinary redemptions at time t_n, while c is the coupon rate on the loan with frq yearly terms. CPR is calculated according to the specified prepayment model, and dependent on the term structure and other variables in the prepayment model. The first factor $\prod_{i=1}^{n-1}(1-CPR_i)$ in (6.3) is the nominal amount of investment left at time t_n. The second factor is the cash flow of a nominal of one at time t_n of the mortgage bond. The product of these two parts gives the cash flow of the bond given an initial investment of nominal one.

Hence, combining (6.1), (6.2) and (6.3) yields the value of the mortgage bond

$$P(0) = \sum_{n=1}^{N} E_0^Q \left[e^{-\int_0^{t_n} r_x dx} \cdot \left(\prod_{i=1}^{n-1} (1 - CPR_i) \right) \cdot \left(CPR_n + (1 - CPR_n) \cdot o_n + \frac{c}{frq} \right) \right] \tag{6.4}$$

Equation (6.4) provides us with the fair value of the mortgage bond. We see that, as we noted in the introduction, in order to calculate the fair value of the mortgage bond, we need both a term structure model and a prepayment model. Like we argued, to discount cash flows, we just need a relevant yield curve (like the one we derived in section 3.1), but in order to make estimates of prepayment rates we need both a term structure model and a prepayment model.

Now that we have theoretical formulas for the value of the mortgage bond, we turn towards how to implement them in the concrete set-up, where we use an interest tree to model the stochastic evolution of the term structure. The interest rate tree provides the basis for calculating prepayment rates and fair values in the model.

The valuation of a callable bond is conducted through backward induction. This method provides us with starting conditions, since the price at the terminal nodes should of course be equal to the cash flow. The terminal node is considered to be an instant before the last payment, which means that the cash flow at the last node consists of the outstanding notional plus the (quarterly) coupon, and is certain at this time. At the nodes in the second-to-last and earlier stages, things get somewhat more complicated. This is due to the stochastic nature of the cash flow at non-terminal nodes. At non-terminal nodes, the CPR should be estimated. Remembering that the cash flow of the mortgage bond at any non-terminal node consists of prepayments, ordinary redemptions and interest payments, we just denote the cash flow at node (i, j) by $CF(i, j)$ and write the value of the mortgage bond at terminal nodes as

$$P(i, N) = CF(i, N) \quad \forall i \tag{6.5}$$

while the value of the bond $P(i, j)$ at non-terminal nodes (i, j) can be written as

$$P(i, j) = CF(i, j) + E_j^Q \left[\{PV(P(k, j+1))\}_{k=i-1}^{i+1} \right] \tag{6.6}$$

The martingale expected value in stage j of the price in stage $j + 1$ is simply calculated using the derived probabilities from section 3.4. Hence, by working backwards through the tree (for $j = N$ to 1), each time calculating $P(i, j)$'s for all i, one obtains fair values for the mortgage bond in all nodes of the tree. Thus, in the end, the value is also obtained in the initial node, which is the aim of the entire exercise.

Note that the pricing method indicates why it is difficult to include path dependent variables in the pricing model. Since the fair value at the initial node is calculated through backward induction, it is difficult to keep track of path dependent variables through the tree. If one finds it essential to include path dependent variables such as the pool factor, it may be a good idea to consider another pricing method, such as Monte Carlo simulation. However, for the purpose at hand and with the prepayment model specified in section 5.3, the pricing method outlined above is satisfactory.

We are now in principle equipped with all the necessities to set up a complete mortgage bond pricing model, since we have now completed both the investigation of how to set up and employ a term structure model and a prepayment model. Furthermore, in this section we have briefly outlined how one in principle could combine these into a pricing model. We therefore conclude the pricing sections here, and in the next sections, we look more at the investment issue, going through return and risk measures that are relevant in the context of callable mortgage bonds in section 7, and then discussing investment strategies in section 8.

7 Return and Risk Measures

Having presented the principles of pricing Danish mortgage bonds, we now present and discuss the most important return and risk measures, which enables us to create hedging strategies as well as general trading strategies in the next section.

Many bond key figures are only sensible to use on plain vanilla bonds and though such key figures have limited theoretical use in mortgage bond analysis, practitioners still find them useful. We will discuss how one adjusts these measures such that they become more reliable measures. We split the presentation of the key figures into two subsections; one for general bond key figures and a section especially for callable mortgage bonds. We end this section by dedicating the last part to the application of some of the presented key figures. This section includes calculations of key figures,[94] such that the reader can get an impression of how these figures capture the characteristics of mortgage bonds.[95]

7.1 General Key Figures

We start out by introducing return figures. For completeness and convenience, we restate the definition of the yield to maturity of a bond. Rewriting (3.1) we have

$$P = \sum_{i=1}^{T} \frac{CF_i}{(1 + \frac{\text{yield}}{frq})^{t_i \cdot frq}} \tag{7.1}$$

where frq denotes the number of payments per annum while CF_i denotes the cash flow that the investor receives at time t_i. Equation (7.1) thus provides us with the annual return of a bond assuming that the bond is held to maturity and that the received cash flows are reinvested at the *same* rate of return. Hence, it is an implicit assumption that the yield curve is flat. Due to this assumption, it is also called the promised yield, as the life of the bond can end prematurely in case of credit events or the bond being called. When calculating the yield of a callable mortgage, one can attempt to correct for the option by using a probability-weighted cash flow from the prepayment model to calculate an expected yield.

The general consensus is that using yield as a decision variable is clearly inferior to NPV measures.[96] Nevertheless, practitioners often use it as a quick-and-dirty indicator of

[94] The key figures are calculated using Danske Analytics, which is kindly made available by Danske Research. For general information on Danske Bank's mortgage bond model, see Danske Research (2002).

[95] In this section we use Danske Research (2004), Hull (2003) and Grinblatt & Titman (2002).

[96] See e.g. the seminal paper Hirshleifer (1958).

profitability just as multiple analysis is used to price equity.

The bond yield is of course not conclusive for our interest in return measures. Our main interest lies with the return we can expect to earn by holding a bond portfolio for a given period of time. In general terms, this is called the holding period return (HPR). HPR consists of both cash flows of the bond – interest payments, instalments and prepayments – and price changes. Both the expected the cash flows and the price change of the bond is dependent on the interest rate expectations. The most conservative approach is to assume that the yield curve is unchanged throughout the period. However, this is only recommended for fairly short investment horizons, while for longer horizons the investor must incorporate his expectations to interest rates, prepayments etc.

To evaluate the attractiveness of a return, one must also evaluate the embedded risk of a strategy. We therefore turn to means of measuring risks of a bond.

The price of a bond paying a fixed coupon or a fixed principal is influenced by (relevant) changes in the yield curve. This is naturally called interest rate risk. The most elementary measure for the interest rate exposure is probably the *basis point value (BPV).* It is defined as minus the first derivative of the price with respect to the interest rate curve

$$BPV = -\frac{\partial P}{\partial r} \tag{7.2}$$

BPV is also called delta risk, as delta risk refers to first derivative with respect to the interest rates. We do not have a functional form for the price-yield relationship, and we must therefore use a numerical approximation using our pricing model to estimate the BPV. We write the approximation as

$$\text{BPV} = \frac{P(r + \Delta r) - P(r - \Delta r)}{2\Delta r} \tag{7.3}$$

where Δr represents the shift in the term structure. In most standard analytics packages a parallel shift is used, but in general nothing prevents the analyst from choosing the shift, which he finds interesting. It is worth noting that due to the different nature of drivers of short and long interest rates, shifts in the term structure are rarely parallel. Short rates are thus primarily driven by monetary policy, while long rates are driven by inflation expectations and supply and demand from the long end investors. Returning to (7.3), $P(\cdot)$ is a model price and it is obtained using a general pricing model as the one in section 2. However, if we use the full mortgage bond pricing model to obtain the bond prices for a callable mortgage bond, it provides us with some degree of adjustment for the embedded prepayment option. The calculated BPV is then called the option-adjusted-BPV. However, for a callable mortgage bond, one is better off using another adjusted key figure, namely ABPV, which we present in the next section.

Duration is another well-known risk measure which is closely related to BPV. The standard duration measure is defined as

$$\text{Duration} = \sum_{i=1}^{N} \left[\frac{PV(CF_i)}{P} \right] t_i \tag{7.4}$$

It can be seen that duration is a weighted average of waiting times for receiving the cash flows, where the weight on each time is proportional to the discounted cash flow at that time. Duration can thus be interpreted as the average time it takes investor to receive the cash flows of the bond. For a zero coupon bond, duration is always time to maturity while for a straight coupon bond, duration is normally be more than half the time to maturity as the principal often constitutes a relatively large share of the cash flows.

By rewriting (7.4), we can obtain another interpretation of duration. It can be shown that duration can be rewritten as

$$\text{Duration} = -\frac{\frac{\partial P}{P}}{\partial \text{yield}} \tag{7.5}$$

However, for us to use duration as an average time for investor to receive the cash flow, we need the yield curve to be flat. Recall that when applying yield one assumes a flat curve. That is, only in case of a flat yield curve can one view duration as minus the percentage change in the price from a change in the yield. The key is to realize that if the yield curve is flat, we can replace ∂r in (7.2) with ∂yield. Hence, rewriting (7.5) provides us with a duration formula based on BPV, which we already have an approximation for. That is,

$$\text{Duration} = \text{BPV}\frac{1}{P} \tag{7.6}$$

Duration and BPV can thus easily be translated into each other. From (7.3) we have an approximation of BPV, which can be calculated from more advanced pricing models. When using such price estimates, one should use the aforementioned interpretations with caution. For now, we just note that the two key figures measure the same interest rate risk and when hedging a portfolio one therefore only needs one of them.

Besides standard duration, there exist a couple variations of the duration measure, which each have their applicability. Macaulay duration denotes the price elasticity with respect to a 1% change in (1 + yield). This is mainly used in simple analytic packages, which can be found in e.g. Microsoft Excel. As it also makes use of the bond yield, it implicitly assumes that the yield curve is flat, which means that it has limited applicability. Another, more applicable, measure is Fischer-Weil duration, which calculates the present value of the cash flow using the zero coupon yield curve instead of the bond yield. Though less dependent on restrictive assumptions, it has a different drawback which is its computational cost. Having an estimated yield curve as the one presented in this thesis, it includes of course no extra calculation. But if this is not the case, the analyst must carry out the exercise we have done in section 3.1.1 in order to obtain a yield curve. When creating investment strategies, an investor might need information on future risks, in which case he would need to model the entire term structure model as we have done it in section 3.4.

Investors often also calculate the exposure towards different key interest rates, that is $\frac{\partial P}{\partial r_i}$, as duration can only be used meaningfully for parallel shifts in the yield curve. We mentioned initially that interest rate risk is referred to as delta risk and a collection of partial derivatives is consequently called a delta vector. The delta vector enables us to decompose the interest rate risk into different rates, which in turn enables us to bet on or hedge non-parallel shifts in the yield curve. Furthermore, the sum of the delta vector entities equals BPV by the property of an integral, and a perfect delta vector hedge thus implies a perfect BPV hedge though the opposite is not necessarily true.

When the interest rate shifts are no longer very, very small, *convexity* becomes important. Convexity, also called gamma, is defined as the second derivative of the price with respect to the yield curve

$$\text{Convexity} = \frac{\partial^2 P}{\partial r^2} = -\frac{\partial \text{BPV}}{\partial r} \tag{7.7}$$

and it thus measures the curvature of the price-yield relationship. As with BPV, we must use a numerical approximation to calculate convexity. We can thus write

$$\text{Convexity} = \frac{P(r + \Delta r) + P(r - \Delta r) - 2P(r)}{(\Delta r)^2} \tag{7.8}$$

For a large absolute value of convexity, the impact of a change in the yield curve becomes inaccurate, when measured using BPV or other first derivative approximations. Therefore, when hedging using BPV, one should always take convexity into consideration. Also, if one applies Macaulay or Fischer-Weil duration, it would be natural to also apply the equivalent convexity measures. Even including the second order derivative of the price-yield relationship is approximative, and in principle one should also apply derivatives of order higher than two. The use of two derivatives corresponds to fitting the price-yield curve locally with a second degree polynomial. However, calculating third or fourth degree derivatives would be to waste a lot of effort on a very small issue. We settle for the first and second derivative.

With a first and second derivative approximation, we have the necessary tools to assess the interest rate risk of a bond portfolio. A plethora of risk measures addresses interest rate risk, but the ones presented here will suffice for our portfolio analysis in section 8. We now turn to risk measures aimed especially at analyzing callable mortgage bonds.

7.2 Key Figures for Callable Mortgage Bonds

In practice, the general key figures presented so far are also used when analyzing callable mortgage bonds, but we need to supplement these with some additional key figures. These enable us to isolate the particular properties of callable mortgage bonds.

When analyzing callable mortgage bonds, one of the most important key figures is *option adjusted spread* (OAS), which is mainly used as a complement to the price quote. OAS is calculated as the spread that must be added to the yield curve for the theoretical price to equal the market price. We write OAS as

$$P = \sum_{i=1}^{N} E^Q[PV(CF_i)]e^{-\text{OAS}\cdot t_i} \tag{7.9}$$

The measure implicitly assumes that the pricing model is correct and it provides a measure for conditions not priced into the bond price by our model. From (7.9) it can be seen that the richness of a bond is inversely related to OAS. A positive OAS can thus be interpreted as a yield equivalent discount for the risk not accounted for by the model e.g. liquidity risk, issuance risk and credit risk. Though it is tempting, one cannot use OAS as an absolute measure for whether a bond is rich or cheap. Instead one must evaluate whether the difference in OAS is a fair value compensation for differences in risk.

In Figure 7.1, we have plotted the price and OAS of RD 5% 2038. Initially the bond trades well above par, and it is therefore closed for issuance. However, the bond price approaches par and hence opening for issuance as interest rates increase. As this will increase the general supply of 5% 2038 bonds, the value decreases and the price thus drops further. The arbitrage-free pricing model does not incorporate supply and demand conditions as this does not influence the fundamental value of the underlying cash flow. This implies that the change in price is mainly reflected in OAS. Interest rates subsequently

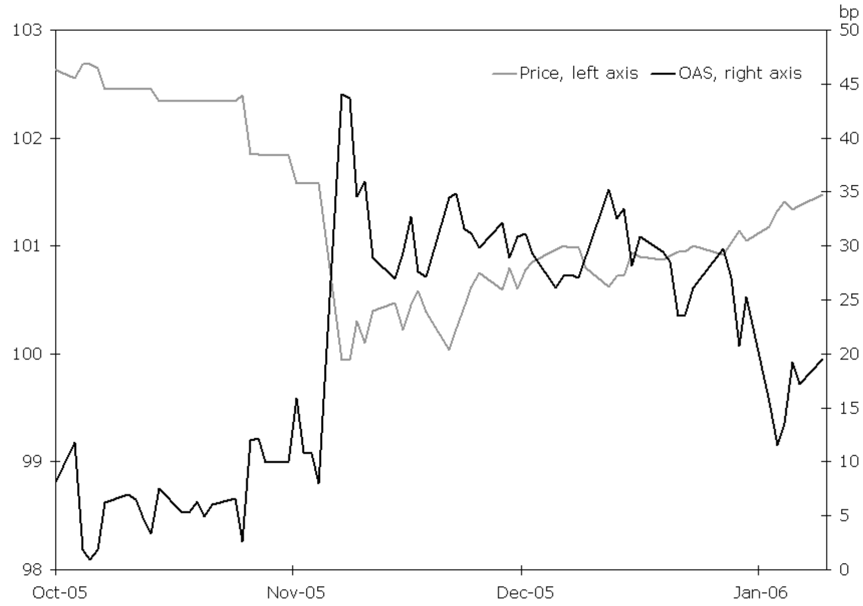

Source: Danske Research

Figure 7.1: Price and OAS, RD 5% 2038

decline again, and the issuance premium of 5% 2038 rightly declines leading to declining OAS. Hence, OAS should be used as a conditional comparative tool to assess if the OAS spread between two bonds is justified by differences in liquidity, credit worthiness, issuance risk etc.

We note that interest rate changes can lead to OAS changes, which in inherently can spill into the price. Hence, the standard BPV measure may estimate interest rate risk incorrectly. In search of a more true estimate, one can adjust BPV with the OAS risk, which is defined as the change in the bond price from a change in OAS. Formally, it can be written as[97]

$$\text{OAS risk} = \frac{\partial P}{\partial OAS} \tag{7.10}$$

Using BPV and ignoring the covariation between OAS and interest rates results in a misleading estimate for the interest rate risk for a callable bond. Hence, we adjust BPV by incorporating the OAS risk accordingly and get the adjusted BPV (ABPV) as

$$ABPV = -\left(\frac{\partial P}{\partial r} + \frac{\partial P}{\partial OAS}\frac{\partial OAS}{\partial r}\right) \tag{7.11}$$

[97]Note also that OAS risk is called spread risk when calculated for a non-callable bond. The reason that practitioners do not adjust standard BPV for spread risk is that little is gained. The spread between non-callable mortgage bonds and government bonds is to a large extent independent of interest rates.

Of the three terms in ABPV we know the two, but we have to estimate the covariation between OAS and the interest rate. This is usually done by a simple linear regression using recent observations of co-movement in yield and OAS.

Finally, we introduce *prepayment risk*, which provides us with the change in the bond price from changes in the estimated CPRs. It is defined as

$$\text{Prepayment risk} = \frac{\partial P}{\partial CPR} \tag{7.12}$$

It provides the investor with a measure for his exposure towards unforeseen prepayments, and it can thus be used to calculate simple sensitivity analyses should one have a reason to believe that the model mispredicts prepayments.

We now move on to an application to of the risk measures to shed light on the most important differences between non-callable bonds and callable mortgage bonds.

7.3 Application: Interest Rate Risk Differences Between Callable and Non-callable Bonds

We illustrate here the difference in interest rate risk created by the call option using callable mortgage bonds and a government bond. We have chosen 4% 2035, 5% 2035 and 6% 2035 from Realkredit Danmark as our mortgage bonds. They are all closed for issuance and the differences between them should therefore mainly be caused by differences in coupons and thus also differences in the option element; that is, how far each option is in- or out-of-the-money. To carry out the analysis, we apply the key figures introduced above.

Figure 7.2 shows a price curve for each bond calculated for parallel shifts in the yield curve. We recognize the close to linear shape of the price-yield curve for a government bond. Being non-callable it obeys the standard inverse relationship for any interest rate level. More interesting are the curves for the mortgage bonds. We notice that the 6% 2035 bond reaches a significantly higher (model-predicted) price than both 4% 2035 and 5% 2035. The reason for this is that the option embedded in 6% 2035 has been far in-the-money for quite some time, and the pool factor is merely 11%. Investors should therefore expect that only modest prepayments will occur in this series – following the discussion in section 4.3.2 – even if interest rates were to drop further. However, 4% 2035 and 5% 2035 both have experienced limited prepayments, and the prepayment model thus predicts fairly large prepayments should their prices reach 102-103.

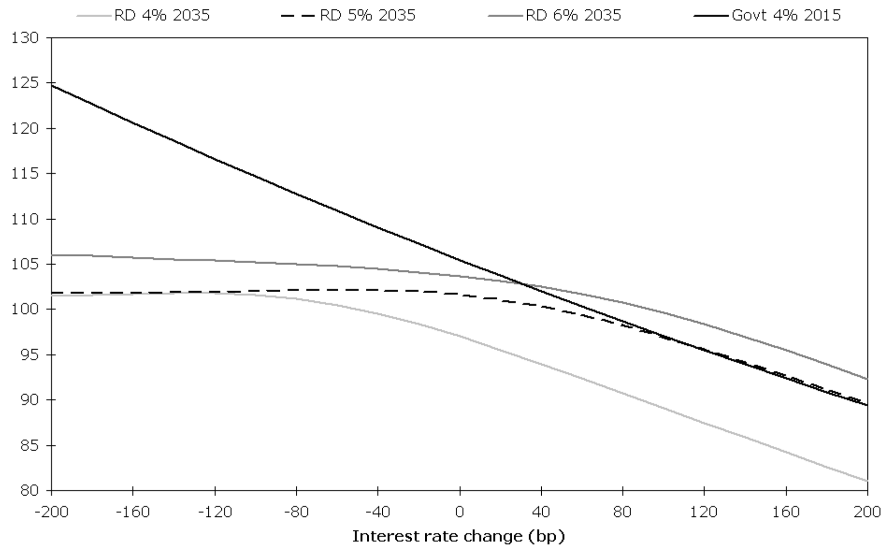

Source: Own calculations conducted in Danske Analytics

Figure 7.2: Price curves for RD 4% 2035, 5% 2035, 6% 2035 and Govt 4% 2015 – January 23, 2006

In Figure 7.3, we have shown the ABPV for the three mortgage bonds and BPV for the government bond. Initially, we note that at the present interest rate levels, the government bond has the highest responsiveness to interest rate changes, while for the mortgage bonds the responsiveness declines with the coupon rate. Also in agreement with Figure 1.1, we see that the government bond has an almost constant decline in price for increases in interest rates. This figure is an excellent illustration of the special

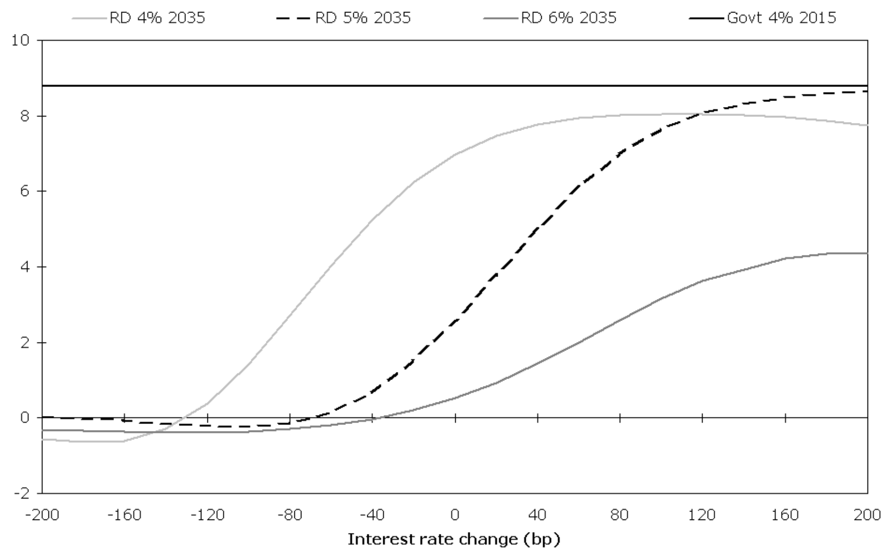

Source: Own calculations conducted in Danske Analytics

Figure 7.3: ABPV for RD 4% 2035, 5% 2035, 6% 2035 and BPV for Govt 4% 2015 – January 23, 2006

characteristics of callable mortgage bonds. For significant interest rate increases, we see that both 4% 2035 and 5% 2035 has roughly the same ABPV as the BPV of Govt 4% 2015 bond. The reason is that the option is far out-of-the-money and the non-callable component is thus dominating regarding the interest rate exposure. As the interest rate decreases and the likelihood of prepayments increases, we see that the option component kicks in. Thus, for sufficiently large interest rate drops, bond prices drop as well, since the prepayment incentive becomes very large. The 6% 2035 ABPV has a somewhat particular pattern. The reason for this is that even if the yield curve was to increase by 200bp, the prepayment option would still be in-the-money. Hence, the ABPV is mainly dictated by the prepayment prediction over the interval, which leads to the particular shape.

Recall that we by calculating the delta vector get the interest rate exposure towards different rates. In Figure 7.4, we see that there is a noticeable difference between the delta vector of a callable mortgage bond and that of a government bond.[98] The government

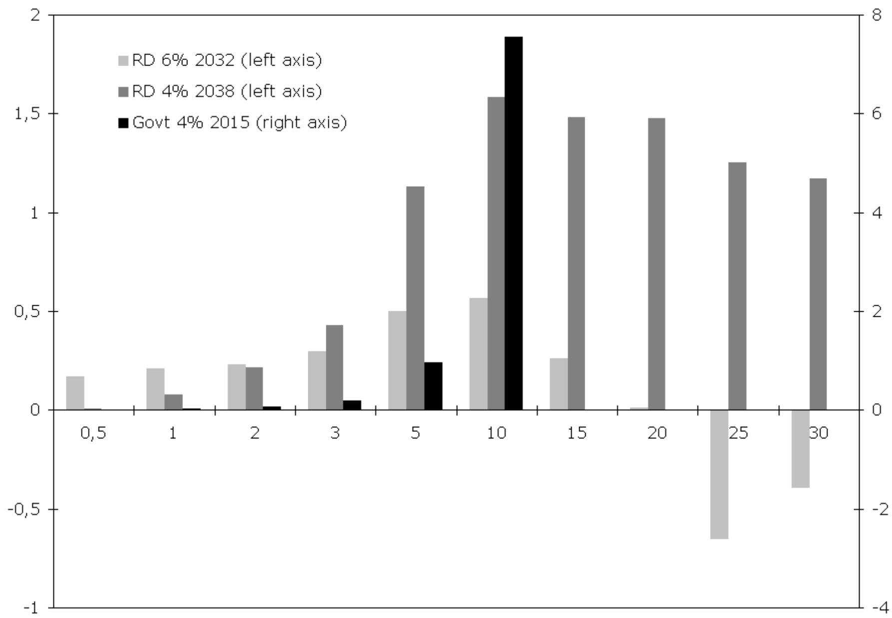

Source: Own calculations conducted in Danske Analytics

Figure 7.4: Delta vector for RD 4% 2038, RD 6% 2032 and Govt 4% 2015 – January 23, 2006

bond is a bullet bond and thus has its main cash flow at maturity. As a result, we see that the coupons result in limited interest rate risk, while most of the risk is affiliated with the 10 year rate. Quite a different picture is seen for the mortgage bonds as they are amortizing callable bonds. 6% 2032 is a premium bond of which the call option is in-the-money. There is thus generally limited potential for price increases, as this would lead to prepayments and a loss of the premium paid. With close to 25 years to maturity,

[98]We have chosen different bonds from those stated in the introduction to this section as they are more suitable for the purpose at hand.

there is a significant negative delta for the 25 year rate as a decrease in this rate pushes the call option further into-the-money.[99] We see a somewhat different pattern for 4% 2038, which is a discount bond and its embedded option is thus out-of-the-money. Hence, yield decreases always lead to price increases and all the delta vector entities are therefore positive. However, unlike the government bond, which is a bullet bond, the 4% 2038 is an annuity bond. As we move along the yield curve two effects influence the interest rate risk. The first effect is the discounting effect, which refers to the fact that a change in the yield curve affect longer horizon cash flows more than short horizon cash flows due to compounding interest. The second effect is called the redemption effect. As time advances, an increasing part of the instalment is made out of redemptions, which limits the interest rate risk for longer horizons compared to shorter horizons. Initially, the discounting effect dominates the redemption effect, but this changes at 10 years where the redemption effect becomes dominating. This results in positive, but declining deltas from 10 year and onwards. To sum it up; premium mortgage bonds can have both positive and negative deltas while discount bonds only have positive deltas.

Now turning to the key element of callable mortgage bonds, we look at the curvature of the price-yield relationship. The property of the curved relationship is called *negative convexity*, though *concavity* would be a more logical term. However, this term has never obtained foothold in the literature. Based on Figure 7.2 and 7.3, we expect 4% 2015 to have close to zero convexity and the mortgage bonds to have interest rate dependent convexity due to the option element. In Figure 7.5, we have plotted the convexity of the bonds for different yield levels using parallel shifts.

The convexity for 4% 2015 is indeed close to zero, but there is a positive small curvature. The negative convexity of the mortgage bonds is highest for the bond which is prone to trigger the highest CPR. In other words, the negative convexity peaks where the downward pressure on the price is highest. Recall from (1.1) that a holder of a callable mortgage bond has shorted a call option. An option has the highest convexity at its strike price, and consequently the mortgage bond has the highest negative convexity around par. Hence, it is no surprise that 5% 2035 has the highest negative convexity. 4% 2035 is a discount bond as mentioned earlier and the call option exerts thus limited downward pressure on its price. Having not had significant prepayments so far, 4% 2035 has the highest potential convexity, but for the current interest rates it is still lower than that of 5% 2035. At the same time, the convexity of 6% 2035 is lower as the prepayment option

[99]Recall, assuming neutral behavior implies that the 25 year rate is used to calculate the refinancing gain.

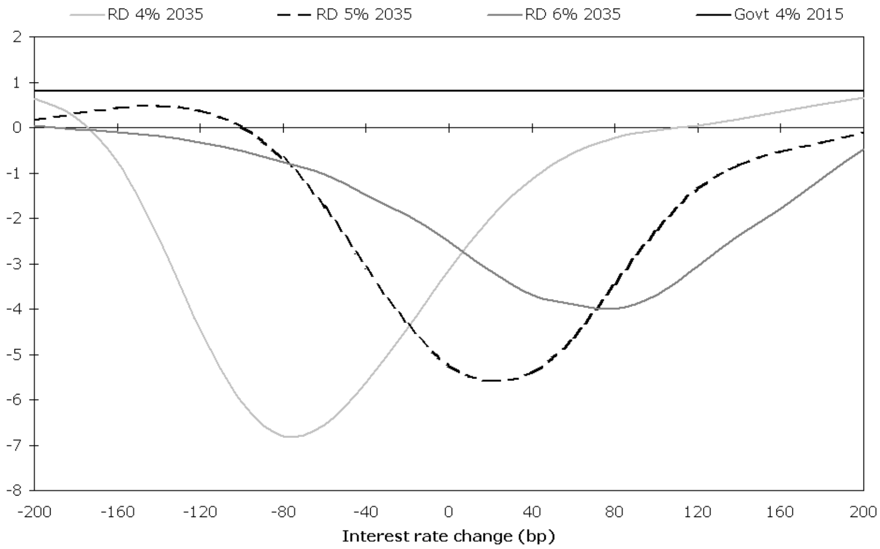

Source: Own calculations conducted in Danske Analytics

Figure 7.5: Convexity for RD 4% 2035, 5% 2035, 6% 2035 and Govt 4% 2015 – January 23, 2006

of 6% 2035 is far in-the-money, since the burn-out of the 6% 2035 bond is significant.

Using the key figures and the knowledge we have obtained on Danish callable mortgage bonds we now move on to trading strategies.

8 Trading Strategies for Mortgage Bonds

An investor can create a portfolio with the intention to either speculate or hedge against market movements. In this section, we investigate how the investor can create such portfolios. In this section we apply the key figures presented in the previous section as well as the knowledge of mortgage bonds that we have obtained throughout this thesis. We start out by presenting how to hedge a portfolio, as this is a vital tool for any investor. After having shown how to hedge risks, we move on to designing strategies in search of mispricings.

8.1 Hedging Strategies

Though the hedging technique presented here can be applied to any kind of risk, we restrict ourselves to looking at interest rate risk.

The premise for our investment is that we have a portfolio, which is created with the intention to speculate against certain market movements. That is, our portfolio unhedged $\mathcal{P}$ has some risk exposures – in this case interest rate risk – that we wish to hedge. Consequently, we create a portfolio Π, which consists of one unit of $\mathcal{P}$ and δ short units of $\mathcal{H}$ – our hedging portfolio – such that we obtain a new, hedged, portfolio

$$\Pi = \mathcal{P} - \delta\mathcal{H} \tag{8.1}$$

The art of hedging is then a matter of choosing δ such that Π has the desired exposure. Using the bond key figures presented in section 7, we can easily calculate the interest rate risk for Π.

Static Delta Hedging

We start out with a static delta hedge as this is a simple representation of a hedging strategy. A perfect delta hedge implies that we choose δ such that the value of our portfolio is invariant for interest rate movements.

We derive the hedge ratio, which implies that Π is (locally) unaffected by any shifts in the yield curve. As we are creating a static hedge, we keep the time unchanged and only change the yield curve

$$d\Pi = \frac{\partial V^{\mathcal{P}}}{\partial r}dr - \delta\frac{\partial V^{\mathcal{H}}}{\partial r}dr \tag{8.2}$$

where V denotes the portfolio value. We set (8.2) equal to zero and obtain

$$\begin{aligned} d\Pi &= 0 \Leftrightarrow \\ \delta &= \frac{BPV^{\mathcal{P}}}{BPV^{\mathcal{H}}} \end{aligned} \tag{8.3}$$

Using (7.6) we can write this as

$$\delta = \frac{Dur^{\mathcal{P}}V^{\mathcal{P}}}{Dur^{\mathcal{H}}V^{\mathcal{H}}} \tag{8.4}$$

Hence, we note that δ is equal to the ratio of BPVs and a value-weighted duration ratio.[100] One should not be deceived by the term perfect hedge, and we again emphasize the fact that BPV only measures the price change following an infinitesimal shift in the yield curve. We also note that even though we implicitly assumed δ to be constant when we differentiated the portfolio value, the hedge ratio varies as the yield curve shifts. Actually, we can see from (7.7) and (8.3) that the change in δ from yield curve changes depends on differences in convexity of the two portfolios, $\mathcal{P}$ and $\mathcal{H}$. Hence, it is natural to combine a delta hedge with a gamma hedge.

Static Gamma Hedge

If our portfolio has significant convexity, the BPV hedge is only effective against infinitesimally small changes in the yield curve. Therefore, if an investor wants an interest rate risk hedge that is fairly robust, he must hedge convexity as well. Recall that a callable mortgage bond is characterized by negative convexity, which means that the bond price decreases more when yield increases than it increases when yield decreases. Hence, it is an unattractive property from the point of view of an investor with a long position in a bond.

By using a second order Taylor approximation and thereby taking convexity into account we have

$$\begin{aligned} d\Pi &= \frac{\partial \Pi}{\partial r}dr + \frac{1}{2}\frac{\partial^2 \Pi}{\partial r^2}(dr)^2 \\ &= -BPV^{\Pi} \cdot dr + \frac{1}{2}\text{Convexity}^{\Pi}(dr)^2 \end{aligned} \tag{8.5}$$

To hedge a portfolio with convexity exposure is only a bit more challenging than hedging a portfolio with only BPV exposure. To see why, we write the expanded version of (8.2)

[100] We use the term BPV in this section. However, when we treat a callable mortgage bond this refers to ABPV.

as

$$d\Pi = \frac{\partial V^{\mathcal{P}}}{\partial r}dr + \frac{1}{2}\frac{\partial^2 V^{\mathcal{P}}}{\partial r^2}(dr)^2 - \delta\left(\frac{\partial V^{\mathcal{H}}}{\partial r}dr + \frac{1}{2}\frac{\partial^2 V^{\mathcal{H}}}{\partial r^2}(dr)^2\right) \tag{8.6}$$

We see that (8.6) only has a solution if the hedging portfolio has the exact same BPV-convexity relationship as the unhedged portfolio. It is very unlikely that we can find a solution with the hedging portfolio consisting of one asset.[101] To emphasize this, we write the hedging portfolio as a collection of assets (or portfolios)

$$d\Pi = \frac{\partial V^{\mathcal{P}}}{\partial r}dr + \frac{1}{2}\frac{\partial^2 V^{\mathcal{P}}}{\partial r^2}(dr)^2 - \sum_{i=1}^{N}\delta_i\left(\frac{\partial V^{\mathcal{H}_i}}{\partial r}dr + \frac{1}{2}\frac{\partial^2 V^{\mathcal{H}_i}}{\partial r^2}(dr)^2\right) \tag{8.7}$$

Hence, by creating a composite portfolio according to (8.7), we obtain a portfolio which is immune to yield curve changes, small as well as fairly large ones.

However, it is important to keep in mind that the hedge ratios presented above only provide static hedging. Hence, they therefore require that the investor actively maintains the proper hedge ratio as time advances, which can be very costly. Due to transaction costs, an investor will not find it profitable to continuously realign the hedge. The recommended approach would therefore be to decide on a threshold, at which the hedge ratio is realigned. This threshold is a subjective matter, which depends on the risk aversion of the investor and the transaction costs. Alternatively, the investor should apply dynamic hedging by deriving hedge ratios without assuming that time is constant. We refer to Taleb (1997) for applications of dynamic hedging. We now move on to creating strategies with the intention of finding mispricings and in this regard, hedging is a very important element.

8.2 Risk Arbitrage – Picking Up Pennies

In this subsection, we propose two trading strategies of our own using mortgage bonds. These strategies are of the kind called risk arbitrage. The oxymoron – risk arbitrage – describes strategies, where the investor aims at locking in an arbitrage profit using only an incomplete hedge. It is an acknowledgement of the fact that it is very difficult to eliminate all risks when trading one portfolio against another. In recent years, investors have experienced numerous examples of investors ignoring risks embedded in strategies that seemingly gave rise to arbitrage profit. The most famous example of this is probably

[101] We disregard the trivial case where the investor buys and sells the same asset for an arbitrage profit.

the case of the Long Term Capital Management[102], which took on exuberant risks using what turned out to be risk arbitrage strategies instead of plain arbitrage strategies. Consequently, risk arbitrage strategies are popularly called *picking up pennies in front of a steam roller* adhering to the large risk accepted for little compensation.

When an investor creates a bet, he has two choices; a directional bet or a relative value bet. The two options are not mutually exclusive, and a strategy can thus be a combination. A directional bet is an investment where the investor expects to make a profit from a change in the state variables, e.g. a steepening yield curve. A relative value bet is based on a belief of the investor that – for unchanged state variables – the market is going to reprice an asset attribute. Investors might lower the required risk premium for holding negative convexity, in which case assets with negative convexity increase in value for unchanged state variables. Relative bets thus include the search for mispricings, and we now turn to how one can create such a bet.

8.2.1 Is 5% 2035 Fairly Priced?

In the world of Danish mortgage bonds, a classic strategy is to pick up extra return via the negative convexity. Recall that all things being equal, negative convexity is an undesirable attribute for a portfolio as this means that our portfolio increases less in value if the yield decreases than it decreases in value if the yield increases. Risk averse investors therefore require a premium for holding a negative convexity portfolio.

In section 7.3, we saw that RD 5% 2035 is a high negative convexity bond, and we wish to investigate whether its embedded premium is fair. This is done by recreating its risk profile using a tracking portfolio and compare the two portfolio returns. Hence, we essentially create a composite portfolio, where we short 5% 2035 and purchase the tracking portfolio or vice versa. Finally, we evaluate the profitability of our portfolio in a 3 month perspective. We choose a fairly short investment horizon as we will assume that the shape of the yield curve remains unchanged. Furthermore, it also decreases the likelihood of our hedge – being static – needs a realignment.

Creating The Hedging Portfolio

We are looking for mispricing and we, therefore, create a tracking portfolio, which tracks both the BPV and convexity profile of 5% 2035 as closely as possible. Instead of merely matching BPV and convexity at the current level, we seek to match the profiles for

[102] Until its fall, LTCM was a very active investor in the Danish mortgage market due to range of possible relative value bets. Lowenstein (2000).

(parallel) shifts in the yield curve. This makes our strategy more robust and thereby mitigates the need for a costly realigning of the hedge later on.

The hedging process continues like this; initially, we determine which asset class to use in our portfolio. Next, we determine the specific assets. Finally, we determine the portfolio weigths. However, choosing the weights is more complicated than it might seem. They are generally determined simultaneously, as changing the weights can influence both BPV and convexity of the tracking portfolio.

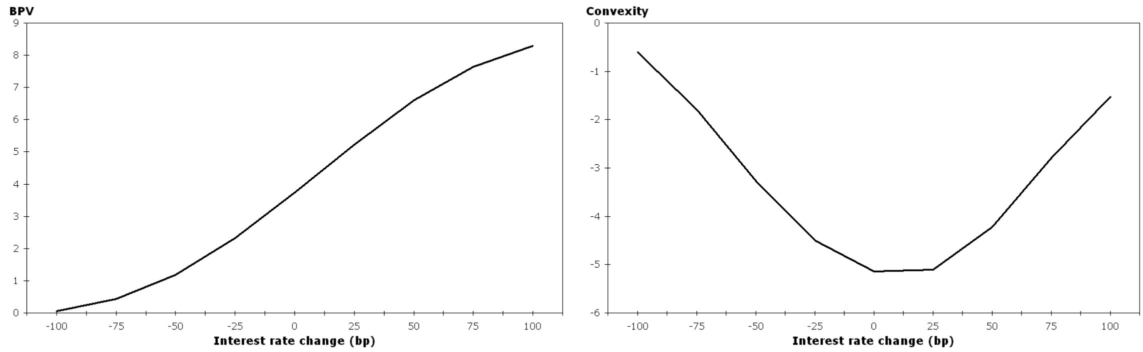

Source: Own calculations conducted in Danske Analytics

Figure 8.1: BPV and convexity for 5% 2035 – February 1, 2006

In Figure 8.1, we have depicted the BPV and convexity for 5% 2035 again.[103] Deciding on which asset types to use for our hedge is fairly simple. Recall from (1.1) that a callable bond consists of a non-callable bond and a short call option. Hence, our tracking portfolio should consist of a bond portfolio and an option portfolio.

To keep the strategy relatively simple, we only use one bond for hedging the non-callable component. The more bonds one allows for, the better match one can obtain, but it also increases the transaction costs. An obvious choice of hedging instrument would be another mortgage bond, as this would enable us to better match delta vectors and thus mitigate curve exposure as well as overall BPV. However, if the premium is mispriced in 5% 2035, maybe it is also mispriced in other Danish callable mortgage bonds. We therefore use a government bond though we are aware of the fact that this implies that we are exposed to changes in the shape of the yield curve. The government bond is chosen to match two factors. First, recall from section 7.3 that a government bond has a fairly constant BPV and convexity unlike 5% 2035. It can thus only be used to level-adjust our hedge and deviations from the constant level of both BPV and convexity should

[103] We limit the interest rate changes of 100bp due to two factors: (1) within the 3 month period in which we evaluate the bet, it seems sufficient, (2) our hedge would be less reliable for a 200bp change in either direction for such a short period of time. Hence, even if we were to create a hedge for a 400bp interval it would be recommendable to realign it, if we experienced larger-sized changes.

subsequently be found in the option portfolio. Second, we choose the government bond such the final strategy is a zero investment – or liquidity neutral – strategy as we otherwise would need to incorporate funding costs.

We use swaptions in our option portfolio, as these provide us with the largest degree of flexibility.[104] DKK Swaption prices are quoted for the interval of options with maturity 1 month to 5 years on swaps with maturity of 1 year to 30 years. Thus we have ample opportunities to create the swaption portfolio we need.[105] To figure out what kind of swaption portfolio we need to create, we again look at Figure 8.1. We can see that 5% 2035 has relatively flat convexity between 0bp and 25bp. Options in general (and thus also swaptions) have the highest convexity level at its strike price and the convexity decreases as the options move in- and out-of-the-money.[106] Hence, we need a short swaption portfolio that mimics the flat level between 0bp and 25bp and furthermore has increasing convexity for drops in the yield curve and increases beyond 25bp. The left leg can be obtained through an at-the-money put option and the right leg can be obtained through an at-the-money+25bp call option. In swaption terms, this means that we use options on a payer swap to obtain the put options and options on a receiver swap to obtain the call options.[107] A strategy consisting of a long put and a long call option, which only differ by the call option having a higher strike price, is called a strangle.[108] Hence, a strangle will enable us to match the convexity profile of 5% 2035.

In Figure 8.2, we show the BPV and convexity characteristics of DKK 1 notional of each swaption.[109] Notice that – prior to the determination of the portfolio weights – taking on these two swaptions seems to enable us to track the convexity exposure of the 5% 2035 profile. The level of convexity can be adjusted by changing the notional amount of the swaptions. However, at this point we only focus on having a similar convexity profile.

The BPV exposure of the strangle is also very satisfactory. Combining the two swaptions gives us an S-shaped BPV risk, which is also recognized from the 5% 2035 profile shown in Figure 8.1. A strangle is BPV-neutral, if it has symmetric strike prices around the current price, as the exposure is then also symmetric around the current price.[110] However, our strangle consists of an ATM and ATM+25bp and we, therefore, expect it

[104] Alternatively, one could construct a portfolio of options on German government bond futures.
[105] See Das (2004) p. 2780 for general information on swaptions.
[106] Hull (2003) p. 315.
[107] The payer leg of an swap pays a predetermined fixed rate while receiving the floating rate e.g. LIBOR+25bp. An option on a payer swap is thus only exercised when the interest rate increases above the predetermined fixed rate. The opposite is true for a receiver swap.
[108] Hull (2003) p. 315.
[109] We have chosen two 6 months options on a 5 years swap.
[110] Duarte, Longstaff & Yu (forthcoming).

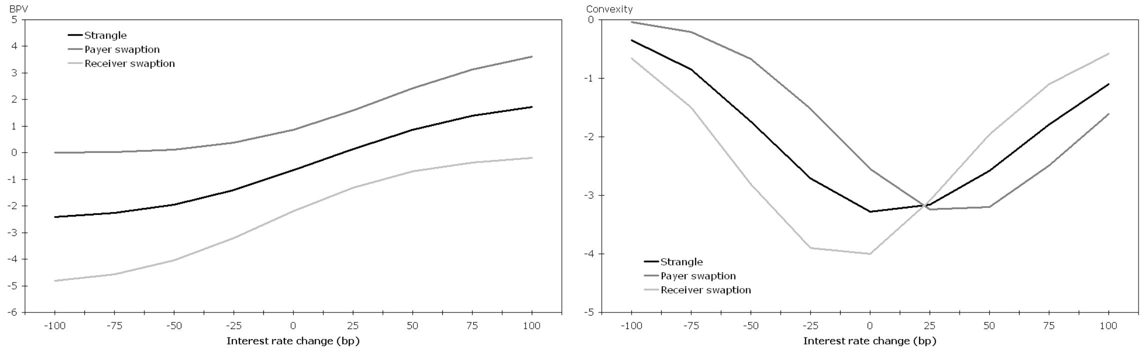

Source: Own calculations conducted in Danske Analytics

Figure 8.2: BPV and convexity for DKK 1 notional of swaption payer and receiver – February 1, 2006

Position	Name	Nominal Amount	Price
Short	5% 2035	-100	-101,78
Long	4% 2010	98.41	103.65
Short	Payer Swaption	-85	-0.18
Short	Receiver Swaption	-80	-0.69

Source: Own calculations conducted in Danske Analytics and price data from Danske Research.

Table 8.1: 5% 2035 and the tracking portfolio – February 1, 2006.

to have positive BPV exposure.

We now determine the weights. One could create a program that determines these by applying e.g. a quadratic loss function using (8.7). However, to do this satisfactorily is a non-trivial matter as it is not obvious how the algorithm should trade-off exposure between states. From a purely theoretical point of view, the deviations should be treated equally, but in practice directionality in the investor's expectations may complicate matters. Say for example that the investor expects interest rates to increase, even in a relative value bet, he is less prone to hedge falling interest rates due to the existence of transaction costs. Treating exposure equally would then be suboptimal. Also, to create such a program is a daunting task that is clearly out of the scope of this thesis. Instead we use a trial-and-error approach.

Looking at the convexity profiles in Figure 8.1 and 8.2, we see that we should increase the notional amount of both swaptions and, at same time, increase the amount of payers relative to receivers. Also, we need to find a government bond, which provides us with a reasonable level-adjustment while turning the combined strategy into a zero investment strategy. The result is shown in Table 8.1 and illustrated in Figure 8.3.

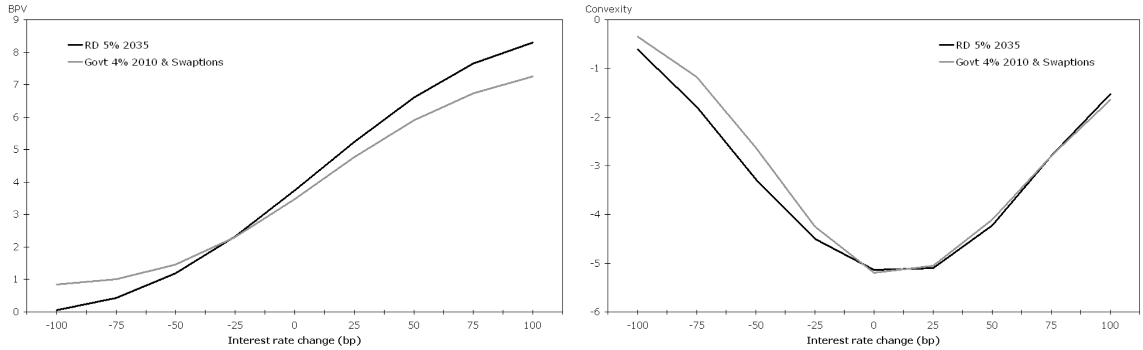

Source: Own calculations conducted in Danske Analytics

Figure 8.3: BPV and convexity for 5% 2035 and hedging portfolio – February 1, 2006

We obtain a highly satisfactory match where we buy notional amount according to DKK 98.41 of 4% 2010, shorting DKK 85 payers and DKK 80 receivers. It can be seen that we obtain an almost perfect match for the convexity profiles though there is a slight deviance for the left leg. This could be improved by adding a payer swaption, which is approximately 10bp out-of-the-money. Turning to the BPV match in Figure 8.3 we note that our match is fairly satisfactory. This could be improved in two ways. Firstly, we could reduce the fit of our convexity and focus more on the BPV hedge using the swaptions we have chosen. Secondly, we should simply short more interest rate exposure for drops in the yield curve and purchase exposure for yield increases. This is done by choosing longer maturities for both the option and the underlying swaps. We do not carry out these computational costly calculations, as we are already highly pleased with our current hedge.

Evaluating the Strategy

To evaluate whether the premium of 5% 2035 is fair, we now look at the return for the two portfolios. From Figure 8.4, we see that the holding period returns for the two portfolios are very close to each other.

Seemingly, we should conclude that 5% 2035 is fairly priced. But recalling the notion of risk arbitrage, we need to investigate the exposures, which we have not accounted for in order to evaluate the bond in total. The most important exposure is the yield curve exposure. As we buy a government bond with 3 years to maturity and short a mortgage bond with 29 years to maturity (or vice versa), we have a mismatch on the exposure to different key interest rates. Consequently, the value of composite portfolio is sensitive towards whether the yield curve change is a steepening or a flattening of the yield curve. However, as a steepening or a flattening can have many shapes and adding the fact that 5% 2035 has negative delta vectors for 25 years and onwards complicates matter excessively.

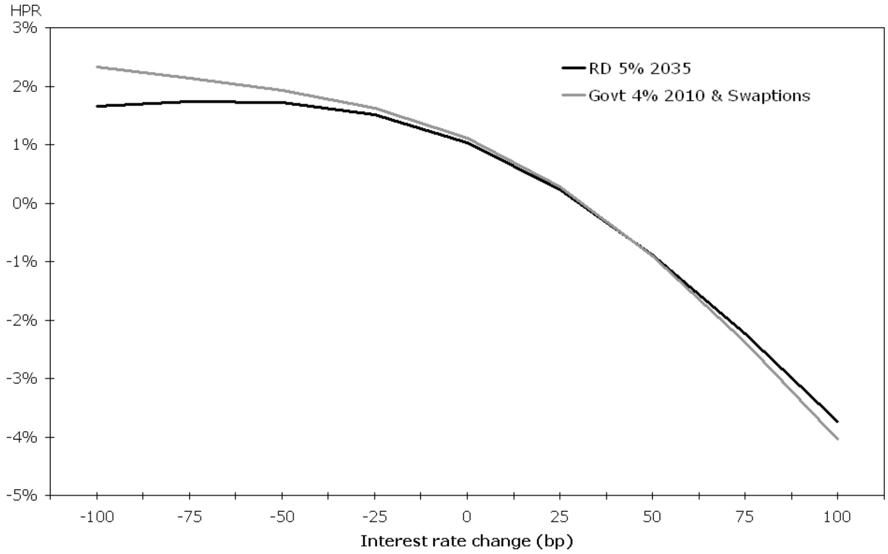

Source: Own calculations conducted in Danske Analytics

Figure 8.4: 3 month HPR for 5% 2035 and tracking portfolio

Hence, in general it is not possible to unambiguously predict the change in value of the composite portfolio for a change in the shape of the curve.

Another non-trivial source of risk, which can change the profitability of our strategy, is a change in volatilities. Anecdotal evidence from the Danish mortgage market indicates that swaptions are more sensitive towards volatility changes than it is the case for the call option embedded in the callable mortgage bond. However, as these exposures do not affect the portfolio return unambiguously in one direction, we are not able to evaluate the fairness of the price of 5% 2035 in general.

To sum up, based on the strategy we have created, we cannot conclude that 5% 2035 is either rich or cheap. Neither buying nor shorting the composite portfolio give rise to a substantial yield pick-up.[111] Adding to this, there are still sources of risk, which we have not accounted for. However, as the impact on the performance of our portfolio is ambiguous, we settle for a conclusion, which states that the interest rate risk from parallel shifts in the yield curve (BPV and convexity) of 5% 2035 is fairly priced.

8.2.2 Prepayment Bet

Another interesting strategy when investing in Danish mortgage bond is a prepayment bet. Based on the statistics on debtor distributions, investors can bet on which issuer gets the highest level of prepayments. A strategy involves shorting a high prepayment bond and

[111] A substantial pick-up means that we across the most probable states are able to create a yield on our tracking portfolio that exceeds that of the portfolio we are tracking or vice versa.

going long in a low prepayment bond as prepayments decrease the value of a bond. Every investor of course takes this into consideration when choosing his general portfolio, which therefore makes this strategy relevant only in the case where the individual manager has a belief that differs from market consensus. Explicit prepayment bets are usually carried out in very young or seasoned series. In the young series, investors have had little chance to observe the behavior of the underlying mortgagors, while the seasoned series having experienced burnout are often difficult to predict as well. To illustrate how such a bet could be constructed, we have constructed the following example.

We look at differences in prices and debtor distributions of the 5% 2038 series. In Table 8.2, we have shown the three largest issues of 5% 2038 bonds. In the table we have shown the debtor distributions of Nordea, Nykredit, and Realkredit Danmark. Let us briefly explain the manner of which the distribution is stated. The distribution for Nykredit is 00-03-24-62-10. This means that respectively 0%, 3%,..., 10% of the outstanding notional amount is parted into loans belonging to respectively DKK 0-0.2 mill., DKK 0.2-0.5 mill., DKK 0.5-1 mill., DKK 1-3 mill. and above DKK 3 mill. We see that the difference in the debtor distributions between Nykredit on the one side and Realkredit Danmark and Nordea on the other is considerable. Nykredit does not only have the largest part above group 3 debtors, but also respectively 2.66 and 4 times as large a share of group 5 debtors as RD and Nordea. As we discussed in section 4.3.3, call options on large loans get in-the-

Name	**Clean Price DKK**	**Debtor Dist.**
NOR 5% 2038	100.725	00-03-24-62-10
NYK 5% 2038	100.725	00-03-11-46-40
RD 5% 2038	100.725	00-04-20-60-15

Table 8.2: Prices and debtor distributions of 5% 2038 – February 14, 2006

money sooner than small loans due to the affine cost structure. Hence, we would expect NYK 5% 2038 to experience higher prepayments as the option comes into-the-money all things being equal. Still, the three bonds trade at the same price, so a seemingly fool proof prepayment bet is to buy RD 5% 2038 (or NOR) and short NYK 5% 2038 in a 1:1 relationship. Notice that such a bet is almost delta and gamma neutral. So, why do prices not differ accordingly? The reason is that it is not fool proof. First of all, debtor distributions cannot perfectly predict prepayments and if debtor distributions are not too different, one should be cautious and not follow debtor distributions blindly.[112] Hence, we

[112]This is a somewhat redundant comment following the rigorous treatment of prepayments presented in sections 4 and 5.

Bond	**Delta**	**Gamma**
RD 5% 2038	3.793	-5.602
NYK 5% 2038	3.730	-5.168
Portfolio	0.063	-0.434

Table 8.3: Prepayment bet for 5% 2038

could see higher prepayments in RD or NOR despite our a priori beliefs about what the debtor distribution suggests. Second, in an increasing interest rate environment, the price of 5% 2038 is more likely to drop and reopen than it is to increase and push the option further into-the-money. In the case where the bond series reopen, the debtor distributions will most likely change, and so will the conditions for the prepayment bet.

This ends the treatment of the investment issue in our thesis. Before we conclude on the findings from the entire thesis in section 10, we briefly discuss the emergence of the new mortgage bond products on the Danish market in the next section, emphasizing the increasing complexity that market participants face.

9 Product innovation

As we explained in the introduction, the Danish mortgage loans have traditionally been long-term callable fixed interest rate loans, and due to the so-called balance principle, the mortgage bonds outstanding have correspondingly also been long-term fixed interest rate bonds.[113] This pattern has changed considerably recent years. The present section describes some of the new innovations on the Danish mortgage credit market in recent years, and briefly discusses the issues involved in pricing the underlying bonds. Hence, this section serves the purpose of illustrating the increasing complexity that market participants face.

9.1 Adjustable Rate Mortgages

The recent development of new products on the Danish mortgage credit market started with the introduction of adjustable rate mortgages introduced by Realkredit Danmark in 1996 under the brand *FlexLån®* (Flex Loans),[114] which became later the commonly accepted name for adjustable rate mortgages on the Danish market. The adjustable rate mortgage loans can be taken on with a list of different maturities, just as is the case for traditional callable mortgages, which means that the maximal maturity for such a loan, according to current Danish legislation, is 30 years.

However, the underlying bonds are very different from traditional long-term callable bonds. One very important issue concerning the adjustable rate mortgages in Denmark is that the maturity of the underlying bonds is shorter than the maturity of the loans. This has to do with the length of the period that the interest rate on the loan is fixed from fixing to fixing. Normally, the fixing period on a Flex Loan is between one and five years, such that the interest rate for the mortgagor is fixed for the next one to five years, respectively.

This is carried out in practice by mortgage banks issuing new bonds every time the interest rate is adjusted. If the fixing period on a Flex Loan is one year, the underlying bonds will also only have a maturity of one year. This means that even though these loans may seem like floating-rate loans to the mortgagors, the underlying bonds are actually fixed interest rate loans. So, in this respect, the bonds underlying the Flex Loans function like the traditional mortgage bonds. Of course, traditional mortgage bonds and Flex

[113] The balance principle is a part of the legislation of mortgage banks in Denmark. It dictates a relatively strict balance between the loan and the funding side of a mortgage loan.

[114] For a short description, consult www.rd.dk.

Loans differ in other important aspects, most importantly the maturity of the bonds, the absence or presence of a prepayment option – the bonds underlying the Flex Loans do not have a prepayment option,[115] and the amortization profile. While traditional mortgage bonds are annuities, the bonds underlying the Flex Loans are bullet bonds. Actually, the mortgagors pay instalments at quarterly terms, but these are not passed on directly to the investors. Rather, the running instalments are collected by the mortgage bank, and passed on to the investors at the maturity of the underlying bonds.[116]

From the mortgagor's point of view, the advantage of taking on a Flex Loan as compared to a traditional long-term fixed interest rate loan is obvious; the interest rate on the Flex Loan will typically be lower than on the long-term fixed interest rate loan, except in cases where the yield curve is inverse (negatively sloped).

The bonds underlying the Flex Loans are relatively easy to price. What complicated the pricing of traditional mortgage bonds was the prepayment option, and for the bonds underlying the Flex Loans, this is not an issue, since these loans are not callable. Therefore, pricing these bonds is simply a matter of discounting the payments on the bond using a relevant yield curve.[117]

9.2 Capped Floating Loans

Even though Flex Loans are attractive for the mortgagor due to an initial interest rate saving in cases of a normally shaped yield curve, it is a risky loan to take, since the payments on this loan could rise unlimitedly in connection with the refinancing auctions. In 2000, the Danish mortgage bank Totalkredit was the first on the Danish market to provide a mortgage hybrid called *Bolig-X lån* consisting of a mortgage loan with adjustable interest rate, but with a cap over the maximal interest rate.

The bonds backing these mortgages are *floaters*. The term *floaters* points to the fact that even though these loans may seem very similar to Flex Loans, the actual construction is very different. While Flex Loans are issued as short-term bullet bonds that are rolled, the floaters are issued as long-term annuity-like bonds, but with an interest rate that is flexible, and adjusted frequently.

Hence, the basic idea is that the interest rate is adjustable, and in the Danish case, it is adjusted every 6 months. The underlying interest rate is a 6-month money market rate

[115] However, the loans can of course be prepaid in connection with the refinancing auctions or by buying the bonds back in the market (exercising the delivery option).

[116] The positive interest of these running instalments are incorporated into the payments on the loan.

[117] If necessary, added or subtracted a relevant spread.

(CIBOR6), to which a spread is added to obtain the coupon rate. However, should the underlying interest rate rise above a predetermined level, the interest rate is capped at exactly this level. Should the underlying interest rate fall again, so will the interest rate on the mortgage. Such a bond is now logically referred to as a *capped floater*.

In the beginning, where only Totalkredit offered these loans, the maturity of the loan and the cap was only five years. So even though the mortgagor could obtain insurance for very adverse interest rate movements, such an insurance could not be obtained for a typical maturity of a Danish mortgage. Jakobsen & Svenstrup (2001) analyzed this loan type at that time, and concluded that this new product on the Danish mortgage market was an important innovation. However, they suggested that the mortgage banks should introduce similar loans with longer maturities. This was based on the observation that the Flex Loans had gained so much popularity in short time, and they concluded that the combination of a low level of interest and insurance for the maximum payments, had to be "the perfect mortgage".

It took a few years for the mortgage banks to follow this suggestion, but in 2004, Realkredit Danmark introduced a 30-year flexible rate mortgage with a cap. However, it was not entirely similar to the capped floaters that Totalkredit had introduced. Instead, the new capped loans were constructed such that the interest rate followed the CIBOR6 interest rate, just as the capped floaters, but in case the interest rate should hit the cap, the mortgage would automatically be converted into a traditional callable fixed-interest rate loan with a coupon equal to the cap rate. In effect, this meant that if the interest rate should fall below the cap rate at a later stage, the interest rate on this new loan type would not change, since the loan had been converted into a traditional callable fixed interest rate loan. These type of loans were later termed *floating-to-fixed*.

Since then, all the Danish mortgage banks have started issuing either capped floaters, floating-to-fixed or both. To emphasize similarities and the differences between these two loan types, we outline the most important differences in Table 9.1.

One of the most difficult problems to tackle when pricing either capped floaters or floating-to-fixed bonds, is the stochastic amortization. On a non-callable fixed interest rate mortgage, the amortization schedule is deterministic. In other words, the size of the payments (both interests and instalments) are known all the way to maturity. The fact that the amortization schedule for the capped floaters of floating-to-fixed is stochastic, stems from the following factors.

– **Prepayments.** The prepayment option on capped floaters (not relevant for the

	Capped Floater (CF)	**Floating-to-Fixed (FF)**
Underlying interest rate	CIBOR6	CIBOR6
Adjusting frequency	6 months	6 months
Interest rate after being capped	Interest rate falls again	Interest rate stays at the cap level
Callability	Yes, at strike 105	Automatically when capped, and the thereby issued bond is callable at par
Instalment-free option	Yes	Yes

Table 9.1: Properties of capped floaters and floating-to-fixed bonds

floating-to-fixed) in itself causes the amortization to be stochastic, in a manner similar to the prepayment issue on traditional long-term fixed interest rate callable bonds.

– **Amortization is dependent on the future interest rate development.** This has to do with the fact that the instalments on a capped floater or a floating-to-fixed are recalculated every time the coupon rate is changed. The amortization schedule is recalculated at every reset date according to an assumption of unchanged interest rates in the future, and on basis of a standard annuity. Hence, the amortization depends on the interest rate in the future. Notice that this is actually an attractive feature of the floaters – at least from the mortgagors' point of view – that the amortization schedule is recalculated every time the coupon rate is changed. An interest rate increase thus causes the initial instalments on the new loan to decline, such that one can say that apart from the cap insurance, there is also some sort of automatic stabilization built into the floaters. Hence, the instalments will be set high in periods of low coupon rates and low in periods of high coupon rates.

– **Emission pattern.** The emission in these series is aggregated by loans that may have different maturities.

When pricing traditional long-term fixed interest rate callable bonds, only the first of these three issues, namely the prepayment issue, is relevant. Hence, this is the only source of stochastic amortization for traditional long-term fixed interest rate callable bonds.

However, when pricing capped floaters (and floating-to-fixed), the most important source of uncertainty (stochastic amortization) is the interest rate development. Hence, this is the factor that usually receives the most attention when modelling capped floaters and floating-to-fixed bonds. The interest rate dependent amortization schedule furthermore causes path dependency in the pricing process. Since there is a (negative) correlation between interest rates and instalments, the payments on a capped floater or a floating-to-fixed, is dependent on what the interest rates have been in the past. This calls for the application of Monte Carlo simulation methods, rather than using the usual lattice approach. Still, when valuing capped floaters, a prepayment model is needed in the Monte Carlo pricing model – a model that builds on simulations of the evolution of the term structure according to the specified interest rate model (e.g. Hull-White). In order not to complicate the model further, a very simple prepayment model is often assumed, e.g. that in case the price should increase to a level above the strike at 105, the fraction of mortgagors that prepay their loans is equal to a constant α.

In the case of a floating-to-fixed, the prepayment option is no longer relevant, since the conversion of the floating rate loan into a fixed interest rate loan is done automatically when the floater is capped at the cap rate. However, since there is a positive probability of the loan being converted into a traditional long-term fixed interest rate callable loan, the value of such a mortgage is needed in all nodes of the interest rate tree to create a fair value of the floating-to-fixed. Therefore, two concurrent pricing models are needed, one for the floater and one for the corresponding fixed interest loan. Monte Carlo simulation is usually used for the floating part of the bond, and this should be combined with an ordinary pricing model for the traditional mortgage, which the floater could eventually be converted into, the calculations of such a model is computationally very demanding, even more demanding than those for a capped floater, which are also very time consuming.

This pinpoints the fact that the product innovation on the Danish mortgage credit market has and will put forward new demands for market participants to be able to apply relatively advanced computational methods in order to price mortgage bonds accurately.

9.3 Instalment-free loans

Since October 2003, it has been possible to take on a mortgage loan, which is instalment-free.[118] In other words, the mortgagors can avoid paying instalments and only pay the interests on their mortgage. The current legislation restricts the instalment-free period to be maximally 10 years, but in effect, the loans can be rolled after 10 years. Hence, in

[118] It was a change in the legislation in June 2003 that made these loans possible. Sometimes these loans are, in the literature, referred to as Interest-Only rather than instalment-free.

effect, it is possible to make a perpetuity. The introduction of instalment-free loans had the potential to dramatically decrease the monthly payments on a mortgage, and consequently the instalment-free loans have gained huge popularity among many real estate owners.[119] According to Danmarks Nationalbank, the instalment-free fraction of mortgage loans to owner-occupied dwellings, amounts to 31.5% of the outstanding amount.[120]

The interest-only option is not restricted to a specific loan type. It is an option to both traditional long-term fixed interest rate loans, Flex Loans, and the floaters with an embedded cap (one form or another).

The simplest case to treat is the instalment-free option embedded in a Flex Loan. Since the bonds underlying the Flex Loans are, as explained in section 9.1, bullet bonds, the way that the instalment-free option is included in a Flex Loan, is simply to adjust the refinancing amount in connection with the periodic refinancing auctions.

In the case of the traditional mortgage bonds, the instalment-free option obviously postpones the amortization of the principal, and therefore the duration of a mortgage bond with an instalment-free option is higher than for a similar mortgage bond that does not embed an instalment-free option.

The way that the 10 years of instalment-free payments are placed in the amortization scheme for a traditional long-term mortgage bond, is different for different mortgage banks. Some mortgage banks only offer the interest-only option as an option on one *consecutive* period, such that the mortgagor can choose an instalment-free period, but only one consecutive period of a length up to 10 years, can be chosen. In some mortgage banks, it is even demanded that the instalment-free period should not only be consecutive, but it should also be located in the beginning of the amortization schedule. Other mortgage banks offer a *clip card* arrangement, where the up to 10 years of instalment-freedom can be placed wherever the mortgagor wishes to, possibly in a non-consecutive manner.

Jakobsen & Svenstrup (2003) analyze the pricing of these bonds. They conduct a scenario analysis of these bonds compared to bonds without the instalment-free option, and they conclude that the price of the bond with the instalment-free option should be around two points lower than the price of a similar bond without this option.

In Figure 9.1, the observed prices of two bonds that are similar, except for the inclusion of the interest-only option in one of them, are plotted against the calculated yield of a

[119]There has been critique of the introduction of these loans from various economists, stating that the sensibility of house owners to price falls in the housing market became very high. The mortgage banks have responded to this critique by resolutely claiming that they would not issue loans to people who would not hypothetically be able to pay a traditional mortgage with instalments.

[120]Danmarks Nationalbank: Statistics on the balance sheets and flows of the MFI sector, December 2005.

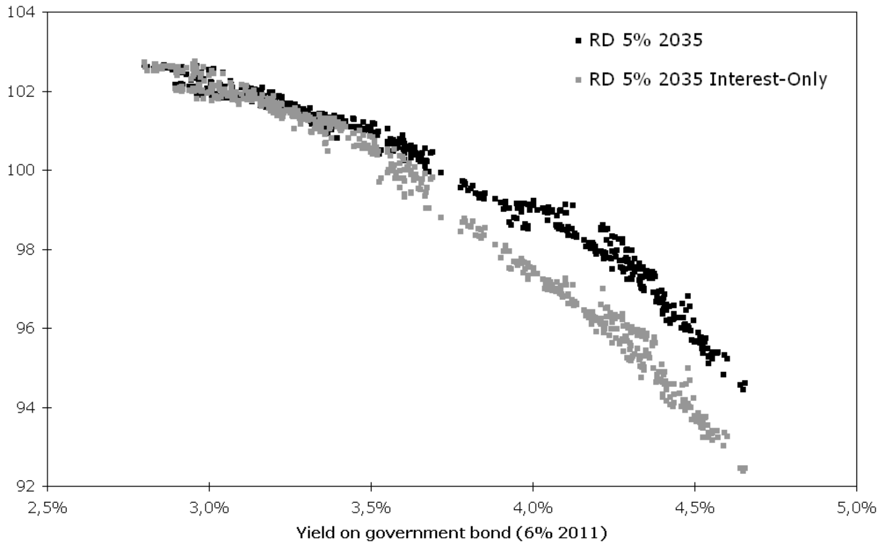

Source: HSH Nordbank Copenhagen Branch

Figure 9.1: Price of RD 5% 2035 with and without the instalment-free option

government benchmark bond (6% 2011). It is seen that the model-based arguments in Jakobsen & Svenstrup (2003) seem to be fairly consistent with the real world observations. For relatively high levels of interest rates, the price difference between the interest-only bond and the ordinary mortgage bond, is approximately two points. Furthermore, it is seen that the price difference narrows when the interest rate falls. This obviously has to do with the prepayment option. As the prepayment option becomes more and more in-the-money, a lot of mortgages will be prepaid. Therefore the significance of the ordinary redemptions[121] declines. When the prepayment option is in-the-money, the important issue is prepayments, since the impact of typical prepayment waves, when the interest rate is below the coupon rate, is most often of a much higher scale. As we have seen, prepayment rates at 15-20% per term are not unusual during waves of mortgage conversions, and in these instances, it is less important whether the ordinary prepayments occur as well.

Another way to look at it is simply that the increasing prepayment threat as the interest rate declines and the bond price exceeds par, causes the duration of both bonds to decline dramatically, and hence, the relative duration difference also declines and the price difference should therefore also decline.

Jakobsen & Svenstrup (2003) furthermore investigate whether it should make a difference that the instalment-free option is given as a clip card or not. They argue that in principle, the mortgagor, by the introduction of the clip card, is provided with an ex-

[121]Which can be calculated using standard annuity formulas.

tra option, namely choosing when to use the instalment-free clips. Therefore, the rational mortgagor would choose to optimize the use of the instalment-free clips, and the bond price with the clip card arrangement should correspondingly be lower than if the instalment-free terms are located as one consecutive period in the beginning of the amortization scheme. However, by conducting calculations in their model, they show that the effect should be minor. In addition to this, it is natural to expect that the primary concern for mortgagors taking instalment-free loans is not rational active debt management, but rather liquidity concerns. This would indicate that even though mortgagors with a clip card would have the opportunity of distributing the instalment-free clips in an optimal way, few of them would probably do so. Instead, the majority of these loans will probably be amortized similarly to the instalment-free loans without the clip card arrangement, such that the full 10 years of instalment-freedom will be placed in the beginning. This fact supports the finding that the price difference between the bonds where the mortgagors have the clip card arrangement and the bonds where they do not, should be minor.

9.4 Future Innovations

As it is seen from the preceding subsections, much has been happening in the Danish mortgage credit market in recent years. Adjustable rate mortgages, capped floaters, floating-to-fixed and instalment free loans are all inventions on the market that have been introduced during the last ten years.

With the product innovation on the mortgage credit market at such a fast pace, many market observers ask themselves what the next innovation will be. It is obviously difficult to tell, but still there may be indications as to what will be logical to introduce for the mortgage banks. If we look at the existing loan palette, the new loans that have gained the highest popularity have all been loans that enabled the mortgagor to pay lower payments on their mortgage. So, there is reason to believe that new products will also be products squeezing the monthly payments on the mortgage down to a minimum. Financing on the short end of the yield curve was the first and obvious suggestion, followed by the interest-only option. The introduction of the capped floaters and floating to fixed makes it natural to expect that the next invention will also be based on some kind of derivative. Here, an interest rate floor would be an obvious suggestion. A floor obviously has a value to the investor, who will be willing to pay extra to get a floor on a floating rate mortgage. This, in turn, ensures lower payments for the mortgagor than with a standard floating rate loan. The mortgagor commits himself to pay a certain minimum interest rate on the loan in case the interest rate should fall below this strike rate. Combined with a capped floater, collars could be introduced. It seems just a matter of time until floating rate mortgages with both floors and collars are introduced on the Danish mortgage credit market. The actual low level of interest rates is a serious blockade to the introduction of floors at the moment, since a floor issued out-of-the-money now, will have very limited value.

If these products are actually introduced, it will be interesting to see if the mortgage banks manage to keep their existing practice of issuing the bonds as hybrids where the floating rate note and the derivative are collected, or if the derivatives will be add-on products to a standard floating rate loan instead.

In any case, the mortgage credit market in Denmark becomes more and more complicated and less and less standardized, and it becomes more and more difficult for market participants to be able to calculate the correct prices on existing fixed income securities, including mortgage products. It will therefore be very interesting to see what the future will bring of inventions.

10 Conclusion

Throughout this thesis we have conducted a thorough investigation of the elements involved in setting up a pricing model for Danish callable mortgage bonds. Hence, we have been focusing on two aspects: (1) Modelling of the term structure, and (2) Prepayment modelling, since these are the two main elements of a pricing model for callable mortgage bonds.

Section 2 provided us with the theoretical background of arbitrage-free asset pricing. We presented the martingale approach, which applies the martingale probability measure, or the Q-measure, to calculate the expected value of a cash flow. Using the well-established result that the expected value under Q is the value of the asset, we derived the asset value dynamics, which is dictated by an arbitrage-free assumption. This provided us with the first main result; the term structure equation for a zero coupon bond. The term structure equation is a partial differential equation, which puts restrictions on bond dynamics. Using the zero coupon dynamics, we derive the price of a derivative on the bond exemplified by an option.

Using the framework from that section, section 3 provided a thorough examination of the complete implementation of a term structure model. We started out by estimating the current yield curve using the Nelson-Siegel method. In the estimation we used a broad selection of bonds issued by Realkredit Danmark, putting a strong emphasis on the practical considerations that one needs to take into account when selecting a sample for yield curve estimation, i.e. the parallel use of callable and non-callable bonds.

To price a non-callable cash flow, we only need the current yield curve, but for the modeler to be able to estimate future prepayments, he must be able to model the evolution of the term structure. To do this, one applies a term structure model. Hence, we subsequently solved and applied the one-factor Hull-White (extended Vasicek) model, which belongs to the arbitrage-free model class. By deriving the asset price formulas using the results from section 2, we were able to calibrate the Hull-White model using Danish and Euro prices for caps and floors. To implement the Hull-White model, we turned to the application of a trinomial interest rate tree. The creation of the interest rate tree served the purpose of illustrating the practical implementation of a stochastic term structure model. In this way, we concluded the first part of the pricing model – the modelling of the term structure.

In section 4, we reviewed the scope for modelling prepayment behavior. Initially, we reviewed rational prepayment behavior, which is based on the assumption that a mortgagor exercises the prepayment option optimally, evaluated solely on the economic gain of doing so. Provided the existence of transaction costs and heterogenous loan sizes, rational prepayment models can describe the varying prepayments that is seen in reality. However, the rational prepayment models are based on too restrictive assumptions to provide a satisfactory description of prepayments, but they do provide us with an overview of the basic incentives behind most prepayment behavior. Hence, the single most important variable in explaining prepayments must definitely be the economic gain of prepayment, and we argued why the intuitive variable $\frac{c}{r}$ could be considered as a good proxy. We furthermore argued that the size of the loan and the maturity of the loan (and the pool factor) could also be expected to be relevant drivers of prepayments.

Next, we turned to the modelling of prepayments. A prepayment model is a model providing an estimate for CPR – the conditional prepayment rate, given a selection of inputs. CPR indicates the share of remaining notional, which is prepaid for that period. We reviewed a model for the American market, which uses $\frac{c}{r}$, maturity, seasonality and the burn-out factor as explanatory variables. The model had been implemented by Goldman Sachs, and Richard & Roll (1989) document that it had a global explanatory power of 95% in the period 1979 to 1988. Following this, we reviewed a Danish prepayment model. FinE, a Danish function library, applies a truncated normal distribution as the basis for the prepayment function. This model includes a net present value gain measure, the burn-out factor, a maturity measure and the slope of the yield curve as well as the change in the yield curve. The FinE model also provides a satisfactory global explanatory power by explaining 86.9% of the prepayments in a sample investigated by Madsen (2005).

Subsequently, we started the treatment of setting up our own prepayment model, armed with the knowledge from the investigation of prepayment drivers, and the investigation of the two commercial prepayment models. We chose to use a probit function, which we estimated through maximum likelihood estimation. We included $\frac{c}{r}$, time to maturity and loan size as explanatory variables in the initial model. The model provided us with significant estimates for the economic gain and time to maturity, while the loan size came out statistically insignificant. In a second model, leaving out loan size and adding the slope of the yield curve as well as the change in the yield curve, we obtained statistically significant estimates for all our explanatory variables, except for time to maturity. Using an adjusted R^2 measure – equivalent to the one that Madsen (2005) uses – our model explained 71.7% of the prepayments in our sample. Finally, we discussed which factors could improve the predictive power of our model. These included the media effect, emergence of new products, further use of the debtor distributions and market expectations. Finally, we estimated final prepayments using preliminary prepayments. We found that a simple power function can capture this effect to some extent.

In section 6, we briefly outlined how to combine the term structure model with the prepayment model to obtain fair values of the callable mortgage bond, based on the interest rate tree. This is a economically fairly simple, but a technically extensive task, and we therefore stick with outlining the principles. The basic idea is to use backward induction to calculate an expected cash flow from the bond under the Q-measure, which gives the value of the bond.

Having presented the principles of pricing a mortgage bond, we turned to return and risk measures in section 7. We presented a selection of measures, which are used to analyze the attractiveness of mortgage bond portfolios. We finished the presentation with an application of a selection of the measures to illustrate the differences in interest rate risk between a non-callable government bond and a callable mortgage bond. We showed that the embedded option gives rise to significant differences in both the first and second derivative with respect to the interest rates. A non-callable bond has a fairly constant first and second derivative, while for a callable mortgage bond these measures depend on how far the option is in- or out-of-the-money. This is one of the reasons why mortgage bond analysis is a highly complex area.

In section 8, we showed how one hedges an undesired risk exposure. Hedging is an intuitively simple discipline. One simply cancels out the undesired exposure using a hedging portfolio, which is created to track the risk characteristics of the asset. To exemplify the principle, we showed how to carry out a delta as well as a gamma hedge. Subsequently, we presented the investment area of risk arbitrage. We set up a portfolio, which replicates the BPV and convexity of a 5% 2035 bond. We found that the tracking portfolio has a very similar 3 month return profile, which led us to conclude that the convexity is fairly priced as we are not able of locking in a convincing risk arbitrage return. We also introduced the notion of a prepayment bet. This is a strategy, which aims at identifying differences in coming prepayments. Such a strategy is mainly relevant in an environment with decreasing interest rates. Alternatively, the interest rate environment should be non-increasing as we would expect few mortgagors would find it optimal to if interest rates were to increase.

We used the final section – section 9 – to supplement our treatment of traditional callable mortgage bonds with a brief discussion of the newly emerged loan types. The discussion also addressed the challenges that the new loan types bring to the pricing model for mortgage bonds.

A Mathematical Appendix

A.1 Derivation of Probabilities

We now wish to derive the probabilities in the interest rate tree given that the branching method is (a). We have the following three equations with three unknowns (p_d, p_m and p_u):

$$p_u \Delta R^* - p_d \Delta R^* = -aj\Delta R^* \Delta t \tag{A.1}$$

$$p_u (\Delta R^*)^2 + p_d (\Delta R^*)^2 = \sigma^2 \Delta t + a^2 j^2 (\Delta R^*)^2 (\Delta t)^2 \tag{A.2}$$

$$p_u + p_m + p_d = 1 \tag{A.3}$$

Remember furthermore that $\Delta R^* = \sigma\sqrt{3\Delta t}$. Plugging this into (A.1) and (A.2) and rearranging gives us:

$$\begin{aligned} p_u - p_d &= -aj\Delta t \\ p_u + p_d &= \frac{1}{3} + a^2 j^2 (\Delta t)^2 \end{aligned}$$

Combining these two equations yields

$$p_d = \frac{1}{3} + a^2 j^2 (\Delta t)^2 - p_d + aj(\Delta t)^2 \Leftrightarrow$$

$$p_d = \frac{1}{6} + \frac{a^2 j^2 (\Delta t)^2 + aj\Delta t}{2} \tag{A.4}$$

From the expression for p_d, p_u and p_m are easily calculated:

$$\begin{aligned} p_u &= p_d - aj\Delta t \\ &= \frac{1}{6} + \frac{a^2 j^2 (\Delta t)^2 - aj\Delta t}{2} \end{aligned} \tag{A.5}$$

$$\begin{aligned} p_m &= 1 - p_u - p_d \\ &= \frac{2}{3} - a^2 j^2 (\Delta t)^2 \end{aligned} \tag{A.6}$$

B Programming Appendix

B.1 Nelson-Siegel Estimation

This section lists the functions needed for doing the estimation of a Nelson-Siegel and Svensson yield curve.

Declaration of public constants

```
   Public Const TinY As Double = 0.00000001
Public weights As Boolean
```

Functon NSS calculates a Nelson-Siegel spot interest rate from the parametrical input.

```
   Function NSS(m As Double, beta0 As Double, beta1 As Double, beta2 As Dou-
ble,
beta3 As Double, tau1 As Double, tau2 As Double, modeltype As Integer) As Dou-
ble
If m = 0 Then m = TinY 'avoid error due to division by 0
If modeltype = 1 Then
NSS = beta0 + beta1 * ((1 - Exp(-m / tau1)) / m * tau1) + beta2 * (((1 - Exp(-
m / tau1)) / m * tau1) - Exp(-m / tau1))
ElseIf modeltype = 2 Then
NSS = beta0 + beta1 * ((1 - Exp(-m / tau1)) / m * tau1) + beta2 * (((1 - Exp(-
m / tau1)) / m * tau1) - Exp(-m / tau1)) + beta3 * (((1 - Exp(-m / tau2)) /
m * tau2) - Exp(-m / tau2))
End If
End Function
```

Function PVNSS calculates the price of a bond in Nelson-Siegel or Svensson from the parametrical input.

```
   Function PVNSS(settle As Date, maturity As Date, coupon As Double, fre-
quency As Integer, beta0 As Double, beta1 As Double, beta2 As Double, beta3
As Double, tau1 As Double, tau2 As Double, modeltype As Integer, comptype As
```

```
Integer, datetype As Integer) As Double
Dim Years As Single, fr As Single Dim i As Single
Dim result As Double, m As Double, r As Double
Dim coupondate As Date
numberofpayments = coupnum(settle, maturity, frequency, datetype)
m = YearsTM(settle, maturity)
r = NSS(m, beta0, beta1, beta2, beta3, tau1, tau2, modeltype)
result = 100 * (1 + coupon / frequency) * DF(m, r, comptype)
m = m - 1 / frequency
Do While m > 0
r = NSS(m, beta0, beta1, beta2, 0, tau1, 1, modeltype)
result = result + 100 * (coupon / frequency) * DF(m, r, 1)
m = m - 1 / frequency
Loop
PVNSS = result
End Function
```

Functon YearsTM calculates years to maturity.

```
   Function YearsTM(settle As Date, maturity As Date) As Double
If maturity <= settle Then
YearsTM = 0
Else
YearsTM = (maturity - settle) / 365
End If
End Function
```

Functon DF calculates the discount factor in Nelson-Siegel or Svensson from the parametrical input.

```
   Function DF(Time As Double, Rate As Double, DiscountingMethod) As Double
If DiscountingMethod = 1 Then
DF = Exp(-Time * Rate)
Else
DF = 1 / (1 + Rate / DiscountingMethod) ^ (Time * DiscountingMethod)
End If
End Function
```

The spreadsheet with the functions implemented to estimate the model can be obtained from the authors upon request.

B.2 Probit estimation in SAS

This section lists the SAS programming needed to estimate the probit prepayment function of section 5.3.

```
   PROC IMPORT OUT= WORK.PPDATA DATAFILE="C:\PATH\Filename.xls" DBMS=EXCEL RE-
PLACE;
SHEET="Sheet1$";
GETNAMES=YES;
MIXED=NO;
SCANTEXT=YES;
USEDATE=YES;
SCANTIME=YES;
RUN;

DATA PPDATA2;
SET PPDATA;
NUMBER = 1;
PP2=PP/100;
IF PP = 0 THEN PP2=0.03;
RUN;

ODS HTML BODY="C:\PATH\Filename.html";

PROC PROBIT DATA=PPDATA2 OUTEST=results_1;
model pp2/number = c_r Avg_loan _size Maturity / DISTRIBUTION=NORMAL
ITPRINT;
OUTPUT OUT=b P=Prob;
RUN;

ODS HTML CLOSE;

PROC REG DATA=b;
MODEL PP2=Prob /NOINT;
PLOT PP2*Prob /CFRAME=LIGR;
run;
quit;
```

References

Amemiya, T. (1981), 'Qualitative Response Models: A Survey', *Journal of Economic Literature* **19, No. 4**, 1483–1536.

BIS (1999), 'Zero-Coupon Yield Curves: Technical Documentation', *Bank of International Settlements* .

BIS (2004), 'BIS Quaterly Review, March 2004', *Bank of International Settlements* .

Björk, T. (1998), *Arbitrage Theory in Continuous Time*, 1st edn, Oxford University.

Björk, T. & Christensen, B. J. (1999), 'Interest Rate Dynamics and Consistent Forward Rate Curves', *Mathematical Finance* **4**, 323–348.

Black, F., Derman, E. & Toy, W. (1990), 'A One-Factor Model of Interest Rates and Its Application to Treasury Bond Options', *Financial Analysts Journal* **1**, 33–39.

Black, F. & Karasinski, P. (1991), 'Bond and Option Pricing when Short Rates are Lognormal', *Financial Analysts Journal* **4**, 52–59.

Black, F. & Scholes, M. (1973), 'The Pricing of Options and Corporate Liabilities', *Journal of Political Economy* **81, no. 3**, 637–654.

Brennan, M. J. & Schwartz, E. S. (1977), 'Savings Bonds, Retractable Bonds and Callable Bonds', *Journal of Financial Economics* **5**, 67–88.

Brigo, D. & Mercurio, F. (2001), *Interest Rate Models - Theory and Practice*, Springer Finance.

Cairns, S. E. (2004), *Interest Rate Models - An Introduction*, Princeton University Press.

Chan, K. C., Karolyi, G. A., Longstaff, F. A. & Sanders, A. B. (1992), 'An Empirical Comparison of Alternative Models of the Short-Term Interest Rate', *Journal of Finance* **47**, 1209–1227.

Christensen, M. (2005), *Obligationsinvestering - teoretiske overvejelser og praktisk anvendelse*, 6th edn, Jurist- og Økonomforbundets Forlag.

Cox, J., Ingersoll, J. E. & Ross, S. A. (1985), 'A Theory of the Term Structure of Interest Rates', *Econometrica* **53**, 385–407.

Cvitanic, J. & Zapatero, F. (2004), *Introduction to the Economics and Mathematics of Financial Markets*, The MIT Press.

Dana, R.-A. & Jeanblanc, M. (2003), *Financial Markets in Continuous Time*, Springer Finance.

Danmarks Nationalbank (2005), 'Statens låntagning og gæld'.

Danske Research (2002), 'Danske Banks realkreditmodel'.

Danske Research (2004), 'Obligationsnøgletal'.

Das, S. (2004), *Swaps/Financial Derivatives*, 3th edn, Wiley Finance.

Dixit, A. K. & Pindyck, R. S. (1993), *Investment under Uncertainty*, Princeton University Press.

Duarte, J., Longstaff, F. A. & Yu, F. (forthcoming), 'Risk and Return in Fixed Income Arbitrage: Nickels in Front of a Steamroller', *Forthcoming in Review of Financial Studies*

May 2005.

Duffie, D. (2001), *Dynamic Asset Pricing Theory*, 3rd edn, Princeton University Press.

Dunn, K. B. & McConnell, J. J. (1981), 'Valuation of GNMA Mortgage-Backed Securities', *Journal of Finance* **36**, 599–616.

Dybvig, P. H. (1997), 'Bond and Bond Option Pricing Based on the Current Term Structure', *Mathematics of Derivative Securities* .

Fabozzi, F. J. (2001), *The Handbook Of Mortgage-Backed Securities*, 5th edn, McGraw-Hill.

Gabrielsen, G., Kousgaard, N. & Milhøj, A. (1999), *Likelihood-teori*, Akademisk Forlag.

Giesecke, K. (2004), 'Credit Risk Modelling and Valuation: An Introduction', *Lecture Notes, Cornell University* .

Grinblatt, M. & Titman, S. (2002), *Financial Markets and Corporate Strategy*, 2nd edn, McGraw-Hill Higher Education.

Hawkins, D. M. (1980), 'A Note on Fitting a Regression Without an Intercept Term', *The American Statistician* **34 no. 4**, 233.

Heath, D., Jarrow, R. & Morton, A. (1992), 'Bond Pricing and the Term Structure of Interest Rates: A New Methodology', *Econometrica* **60 no. 1**, 77–105.

Hirshleifer, J. (1958), 'On the Theory of Optimal Investment Decision', *The Journal of Political Economy* **6 no. 4**, 329–352.

Ho, T. S. Y. & Lee, S. B. (1986), 'Term Structure Movements and Pricing Interest Rates Contingent Claims', *Journal of Finance* **41**, 1011–1029.

Hull, J. C. (2000), *Options, Futures, and Other Derivatives*, 4th edn, Prentice Hall International.

Hull, J. C. (2003), *Options, Futures, and Other Derivatives*, 5th edn, Prentice Hall International.

Hull, J. & White, A. (1990*a*), 'Pricing Interest Rate Derivatives Securities', *Review of Financial Studies* **3 no. 4**, 573–592.

Hull, J. & White, A. (1990*b*), 'Valuing Derivative Securities Using the Explicit Finite Difference Models', *Journal of Financial and Quantitative Analysis* **25**, 79–83.

Hull, J. & White, A. (1994), 'Numerical Procedures for Implementing Term Structure Models I: Single Factor Models', *The Journal of Derivatives* **Fall**, 7–16.

Hull, J. & White, A. (1996), 'Using Hull-White Interest-Rate Trees', *The Journal of Derivatives* **Winter**, 26–36.

Jakobsen, S. (1992), 'Prepayment and the Valutation of Danish Mortgage-Backed Bonds', *Ph.D.-Thesis, Aarhus School of Business* .

Jakobsen, S. & Svenstrup, M. (1999), 'Hvad praktikere bør vide om modeller for konverteringsadfærd', *Finans/Invest* **7/99**, 5–11.

Jakobsen, S. & Svenstrup, M. (2001), 'Variabel rente med loft – det perfekte realkreditlån', *Finans/Invest* **07/01**, 18–25.

Jakobsen, S. & Svenstrup, M. (2003), 'Prisdannelsen på afdragsfrie obligationer', *Finans/Invest*

07/03, 9–12.

Jamshidian, F. (1988), 'The One-Factor Gaussian Interest Rate Model: Theory and Implementation', *Working Paper, Merrill Lynch Capital Markets* .

Johnston, J. & DiNardo, J. (1997), *Econometric Methods*, 4th edn, McGraw Hill.

Kvålseth, T. O. (1985), 'Cautionary Note About R^2', *The American Statistician* **34, No. 4 (Part 1)**, 279–285.

Lowenstein, R. (2000), *When Genius Failed; The Rise and Fall of Long Term Capital Management*, 1st edn, Random House.

Madsen, C. (2005), 'Danish Mortgage Modeling in FinE', *FinE Analytics (www.fineanalytics.com) Working Paper* .

Merton, R. C. (1973), 'An Intertemporal Capital Asset Pricing Model', *Econometrica* **41**, 867–887.

Miles, D. (2003), 'The UK Mortgage Market: Taking a Longer-Term View', *HM Treasury Working Paper* .

Moody's (2005), 'Banking System Outlook – Denmark (December 2005)'.

Nelson, C. R. & Siegel, A. F. (1987), 'Parsimonious Modelling of Yield Curves', *The Journal of Business* **60**, 473–489.

Realkreditrådet (2005), 'Realkredit Finansiering i Danmark'.

Rendleman, R. & Bartter, B. (1980), 'The Pricing of Options on Debt Securities', *Journal of Financial and Quantitative Analysis* **15**, 11–24.

Richard, S. F. & Roll, R. (1989), 'Prepayments on fixed-rate mortgage-backed securities', *Journal of Portfolio Management* **15**, 73–82.

Ruppert, D. (2004), *Statistics and Finance - An Introduction*, Springer.

Shimko, D. (2004), *Credit Risk: Models and Management*, 2nd edn, Riskbooks, London.

Shreve, S. E. (2004), *Stochastic Calculus for Finance II - Continuous-Time Models*, Springer Finance.

Stanton, R. (1995), 'Rational Prepayment and the Valuation of Mortgage-Backed Securities', *Review of Financial Studies* **8**, 677–708.

Svensson, L. E. O. (1994), 'Estimating and Interpreting Forward Interest Rates: Sweden 1992-1994', *NBER Working Paper* **4871**.

Sydsæter, K. (2000), *Matematisk Analyse, Bind 1*, 7th edn, Gyldendal Akademisk.

Taleb, N. N. (1997), *Dynamic Hedging: Managing Vanilla and Exotic Options*, Wiley Finance.

Vasicek, O. (1977), 'An Equilibrium Characterization of the Term Structure', *Journal of Financial Economics* **5**, 177–188.

Verbeek, M. (2004), *A Guide to Modern Econometrics*, 2nd edn, Wiley & Sons.

Williams, D. (1991), *Probability with Martingales*, 7th edn, Cambridge Mathematical Textbooks.