Model Risk and Model Choice in the Case of Barrier Options

Rainer Baule[1] and David Shkel[2]

January 14, 2018

[1]University of Hagen, Universitätsstraße 41, 58084 Hagen, Germany, rainer.baule@fernuni-hagen.de

[2]University of Hagen, Universitätsstraße 41, 58084 Hagen, Germany, david.shkel@fernuni-hagen.de.

Abstract

We analyze model risk for the pricing of barrier options. In contrast to existing literature, this paper is based on an empirical data set of over 40,000 bonus certificates to analyze the real market extent of model risk for traded barrier options instead of purely synthetic options. For this purpose a local volatility model, the Heston model and the Bates model are applied. Furthermore, we add to the literature on the behavior of issuers of retail derivatives in terms of model choice. We find evidence that the majority of the issuers prefer stochastic volatility over local volatility models, while they do not use the even more realistic Bates model which incorporates jumps in the underlying.

EFM Classification Codes: 410, 440, 450

Keywords: model risk, Heston model, Bates model, local volatility, bonus certificates, barrier options, empirical finance, retail derivatives

1 Introduction

Derivative pricing is subject to model risk, as the stochastics of the underlying price process and the corresponding market price(s) of risk are unobservable. Besides the usually less important modelling of interest rates, models can broadly be classified into pure-diffusion models (extending the seminal paper of Black and Scholes (1973), which assumes geometric Brownian motion) and jump-diffusion models (introduced by Merton, 1976), and into models with deterministic volatility (such as the local volatility approach pioneered by Breeden and Litzenberger (1978) and further developed by Dupire (1994) and Derman and Kani, 1994) and those with stochastic volatility (one of the most popular being the approach of Heston, 1993). Regarding plain vanilla options, a large number of empirical tests have been presented in the literature, for example by Bakshi et al. (1997) to name one of the most influential studies.

A second-order problem is the pricing of exotic derivatives when the prices of plain-vanilla options are given. In this regard, a trader of exotic options faces *model risk* as the risk that different models yield different prices for the same exotic product even though the models are calibrated to the same observable prices of plain-vanilla options. This paper analyzes model risk for actually traded barrier options embedded in retail derivatives. Furthermore, we provide insight in the preferences of market participants regarding their model choice.

Recent literature has dealt with this sort of model risk from a theoretical perspective, analyzing price deviations for several synthetical exotic options. Hull and Suo (2002) analyze the model risk of compound options and barrier options when traders apply a "practitioner's Black-Scholes model", which is basically a local-volatility model, in comparison to a stochastic volatility model similar to Heston (1993) with fictitious parameters. Hirsa et al. (2003) calibrate the constant elasticity of variance model (Cox and Ross, 1976), two variants of the variance gamma model (Madan et al., 1998), and the local volatility model to S&P 500 index options and compare the model prices of synthetic up-and-out call options. Similarly, Schoutens et al. (2004) calibrate the models of Heston (1993), Bates (1996) (including both stochastic volatility and jumps), Barndorff-Nielsen and Shephard (2001), and some Lévy models with stochastic time to a set of plain vanilla options on the Eurostoxx 50 and analyze price deviations for several sorts of exotic derivatives. Another sort of model risk, namely the dependency of the model outcome on different calibrations, is analyzed by Guillaume and Schoutens (2012).

Besides the impact on derivatives prices, a number of studies also address the sensitivity of hedging performance with respect to different models. Model risk in terms of hedging performance is also an issue for plain vanilla options, as it does depend on the dynamics of the underlying even if market prices of plain-vanilla options are perfectly matched for one instance of time. An influential paper in this regard is Dumas et al. (1998), who find no significant improvement for several calibrations of a local-volatility model over the simple Black-Scholes model. In particular for exotic options, the hedging performance with respect to model choice is addressed by An and Suo (2009), Engelmann et al. (2009), and Nalholm and Poulsen (2006).

All these papers analyze the extent of model risk not for real, but for synthetic options, which is not surprising, as usually exotic options are not traded at public exchanges. Within the segment of structured retail derivatives, there are however products with embedded exotic options, in particular barrier options, with observable market prices. Our paper makes use of these products to present a first comprehensive study of model risk for actually traded exotic options. To be more precise, we analyze a data set of bonus certificates, a product which can be seen as a combination of the underlying with a down-and-out put option. These types of products have already been studied by Baule and Tallau (2011) with regard to fair values obtained by different models. However, they focus on the calibration of the simple Black-Scholes model and analyze model risk with respect to the stochastic volatility approach of Heston as the only alternative. In our paper, we study three substantially different models beyond Black-Scholes: the Heston (1993) model as a representative of the class of diffusions with stochastic volatility, the local-volatility model according to Dupire (1994) as a deterministic approach, and the Bates (1996) model as an approach with jumps in the underlying. As all these models are reasonable and nonetheless different choices for a trader (or, issuer) of bonus certificates, we can quantify the model risk for this product class with an empirical data set of traded products.

In contrast to liquidly traded plain-vanilla options, bid and ask prices for retail derivatives such as bonus certificates are dominated by the issuing bank as a market maker (e.g., Baule, 2011). Thus, observed market prices reflect not necessarily fair values, but are the result of the pricing policy of the respective issuer. This pricing policy has been studied in the literature in a number of directions, including life cycle and order flow (Baule, 2011), competition (Schertler, 2016), and complexity (Henderson and Pearson, 2011). As a second major contribution, we extend this strand of literature by an analysis of model choice.

Both at issuance and in the secondary market, issuers of retail derivatives face the problem of determining a fair value for the product and quoting a price, which consists of the fair theoretical value plus a margin. In this regard, the market for retail derivatives is different from most other financial markets (e.g., Entrop et al., 2016). For plain-vanilla products, such as discount certificates on liquid underlyings, calculating a fair theoretical price is quite easy, as the product is perfectly duplicated by a position in the underlying and a plain-vanilla option, for which a market price at the options exchange exists. Hence, for those products, there is no considerable model risk, which gives issuers the chance to price these products quite aggressively, resulting in low margins (e.g., Baule (2011), Schertler, 2016).

For more complex products with exotic components such as barrier options or multiple underlyings, margins have been found to be considerably larger (e.g., Wallmeier and Diethelm (2009), Stoimenov and Wilkens (2005), Baule and Tallau (2011), Henderson and Pearson, 2011). Reasons for higher margins are an increased opaqueness for investors leading to a greater leeway for issuers to impose a margin, and higher structuring costs. As a part of the structuring costs argument, in contrast to plain-vanilla products, issuers do face model risk in the calculation of a theoretical fair value, since there is no reference exchange for liquid exotic options at which market prices could be observed. Hence, valuation requires the choice of a valuation model and is subject to model risk.

The local-volatility model is often criticized for its unrealistic assumption of a deterministic volatility surface (e.g., Hagan et al., 2002). Traders nonetheless like the model because it has the appealing property that all market prices of plain-vanilla options are perfectly matched (e.g., Hull and Suo, 2002). In fact, traders need not to “believe” in the model and its deterministic volatility, but can use it to price exotic options consistently with plain-vanilla options—similar to the use of the simple Black-Scholes model with adjusted implied volatility to price plain-vanilla options similar to those observed at the market. As Engelmann et al. (2009) show, the local-volatility approach performs satisfactorily for the hedging of barrier options in many cases. In contrast, stochastic volatility models are assumably more realistic, but due to a limited number of parameters not able to capture all plain-vanilla options simultaneously. A first question therefore is:

> Do issuers prefer models with deterministic or with stochastic volatility?

Even more realistic, but also more complicated are models with jumps. The second question to be addressed is:

> Do issuers prefer pure-diffusion models or models with jumps?

Furthermore, there might be differences in the model choice between issuers. A third question therefore is:

> Do issuers agree in their model preferences, or does the model choice vary between different issuers?

The remainder of the paper is organized as follows. Section 2 outlines the construction of bonus certificates and discusses the different valuation models together with an approach to measure model risk. Section 3 presents the methods applied for model calibration. Section 4 presents the empirical data. Section 5 shows the results of the model risk analysis. Section 6 discusses the model choice analysis and answers the questions raised. Section 7 concludes.

2 Valuation Models

2.1 Bonus Certificates

A classic bonus certificate is a combination of two parts. It consists of the underlying (or, a zero-strike call on the underlying) and a down-and-out put option on the same underlying. This construction leads to two different payoff possibilities. In the first case, the underlying never crosses the barrier during the product's lifetime, which results in a payment equal to the value of the underlying at maturity, yet at least equal to a bonus level, which equals the strike price of the down-and-out put. In the second case, the barrier is touched or crossed by the underlying at least once during the product's lifetime, which results in a payment equal to the price of the underlying at maturity. Hence, a breaching of the barrier knocks out the chance of a bonus payment. If S_t is the underlying price at time $t \in [0;T]$, where T is the maturity date, H stands for the barrier and B defines the bonus level, the payoff BC at maturity can be summed up as follows:

$$BC = \begin{cases} S_T & \text{if } \min\limits_{t \in [0;T]} S_t \leq H, \\ \max\{S_T; B\} & \text{else.} \end{cases} \tag{1}$$

An extension of this construction is the so called capped bonus certificate. In this variant, the maximum possible payoff is capped at a certain level. As a consequence, investors do not have the chance of unlimited gains. This additional feature is constructed by including a short call option in the duplicating portfolio with the strike equal to the cap level. With the aforementioned symbols and C describing the cap level, the payoff CBC at time T can be written as

$$CBC = \begin{cases} \min\{S_T; C\} & \text{if } \min\limits_{t \in [0;T]} S_t \leq H, \\ \min\{\max\{S_T; B\}; C\} & \text{else.} \end{cases} \tag{2}$$

To receive the fair value of a bonus certificate or a capped bonus certificate, it is sufficient to calculate and sum up the fair values of all parts of the duplicating portfolio. The fair value of a zero-strike call equals the underlying price minus the present value of dividends paid during the product's lifetime. Due to the fact that this paper examines only options written on the DAX performance index, it is not necessary to incorporate possible dividend payments. Therefore, it remains to calculate the fair value of the down-and-out put to determine the bonus certificate's fair value. To get the capped bonus certificate's fair value, one has to additionally calculate the short call's fair value.

The analysis so far results in a fair value without the consideration of the issuer default risk. In general, securitized retail derivatives are uncollateralized instruments subject to issuer default risk (see Baule et al. (2008) for a deeper discussion). Throughout this paper, the default-free prices are adjusted by using the Hull and White (1995) model to account for default risk. Assuming independence between market risk and default risk, the fair value without credit risk, FV, of a certificate with maturity T, is discounted with the issuer's credit spread c. Therefore, the adjusted value $\overline{FV}$ of the certificate is given by

$$\overline{FV} = e^{-c \cdot T} \cdot FV. \tag{3}$$

2.2 Local Volatility

The local volatility model is based on the work of Dupire (1994) and Derman and Kani (1994). The volatility varies over time and depends on a combination of the given point in time and the corresponding price of the underlying according to a deterministic function. The evolution of the underlying price process is given by the following partial differential equation:

$$\frac{\mathrm{d}S_t}{S_t} = r\mathrm{d}t + \sigma(S_t, t)\mathrm{d}W_t, \tag{4}$$

where S_t denotes the price of the underlying at time t, r is the continuous risk free interest rate, W_t is a standard Wiener process, and $\sigma(S_t, t)$ is the local volatility function.
An advantage of the local volatility model is the possibility to perfectly calibrate it to a given data set of plain-vanilla options. This is possible, since the model has an infinite number of parameters regarding the time dimension and the underlying price dimension. Another advantage is the straight forward simulation of the underlying price process since there exist only one source of randomness in the model whereas stochastic volatility models use at least two sources of randomness. The local volatility model is skew-consistent and features all empirical conditions regarding volatility surfaces due to the fact that a perfect calibration is possible. Therefore, it is a useful tool to price exotic options even though it lacks the dynamics of a real world volatility process.

2.3 Heston Model

The Heston model is a stochastic volatility model. It consists of two correlated partial differential equations, one describing the underlying price process and the other one describing the volatility process. The model is given by the following equations:

$$\mathrm{d}S_t = rS_t\mathrm{d}t + \sqrt{v_t}S_t\mathrm{d}W_t^1, \tag{5}$$

$$\mathrm{d}v_t = \kappa(\theta - v_t)\mathrm{d}t + \xi\sqrt{v_t}\mathrm{d}W_t^2, \tag{6}$$

where S_t denotes the price of the underlying at time t, r is the continuous risk free interest rate, v_t is the variance at time t, W_t^1 and W_t^2 are two correlated standard Wiener processes,

κ is the mean reversion speed, θ the mean reversion level and ξ the volatility of volatility. Furthermore, the initial variance is given by v_0. The two Wiener processes are correlated with the correlation coefficient ρ, that is $\langle dW_t^1, dW_t^2 \rangle = \rho dt$.

The volatility follows a mean reversion process, so the instantaneous volatility tends to evolve in the direction of the long term average volatility level. The higher the mean reversion speed κ, the faster the return to the mean reversion level. By using the Heston model one has to take care of the volatility process because volatility always has to be non-negative. This property is guaranteed if the model parameters fulfill the Feller (1951) condition $2 \cdot \kappa\theta - \xi \geq 0$. This condition prevents the volatility process from reaching zero and therefore from becoming negative since the initial value should always be positive.

The small number of five model parameters leads to a non-perfect model calibration. The result of this imperfect calibration is a good fit for long maturities and some differences between observed and theoretical option prices for short maturities. While the model is able to capture the volatility dynamics for the long haul, it is not able to reproduce the steep volatility smile for short maturities due to the lack of jumps in the underlying price process and/or the volatility process. Typically, the correlation parameter ρ takes negative values. As a result, the volatility tends to fall if the underlying rises and vice versa, which is a well-established empirical fact.

A useful property of the Heston model, which holds also true for the Bates model and a variety of other stochastic volatility models, is the availability of closed-form solutions for plain vanilla options (see Bakshi and Madan (2000) for the pricing formula for an European call option). This is a real advantage regarding model calibration. Unfortunately, there exist no formulas for barrier options, which causes the necessity to use a numerical approach.

2.4 Bates Model

The Bates (1996) model combines the stochastic volatility model of Heston with log-normally distributed jumps in the underlying price process. These jumps occur randomly and are independent of the evolution of the diffusion process and independent of the evolution of the volatility process. By incorporating random jumps, the Bates model features another empirical fact of stock markets, which has been shown for example by

Andersen et al. (2002) or Eraker (2004). The Bates model is given by the two following partial differential equations:

$$\mathrm{d}S_t = (r - \lambda\mu_J)S_t\mathrm{d}t + \sqrt{v_t}S_t\mathrm{d}W_t^1 + J_tS_t\mathrm{d}N_t \quad (7)$$

$$\mathrm{d}v_t = \kappa(\theta - v_t)\mathrm{d}t + \xi\sqrt{v_t}\mathrm{d}W_t^2. \quad (8)$$

The volatility process equals the one from the Heston model. N_t is a Poisson process with intensity λ and J_t is the random jump size which is log-normally distributed:

$$\log(1 + J_t) \sim \left(\log(1 + \mu_J) - \frac{v_J^2}{2}; v_J\right), \quad (9)$$

where μ_J is the expected jump size and v_J is the standard deviation of the jump size. Due to the existence of jumps, one has to adjust the drift rate of the diffusion term to make sure that the expected return in the risk neutral world equals the risk free interest rate. As in the Heston model, the underlying price process and the volatility process are correlated with $\langle \mathrm{d}W_t^1, \mathrm{d}W_t^2 \rangle = \rho\mathrm{d}t$. The negative dependence between the underlying price process and the volatility process is even more pronounced in the Bates model compared to the Heston model. The jumps are independent of the other sources of randomness, that is, $\langle \mathrm{d}W_t^1, \mathrm{d}N_t \rangle = 0$ and $\langle \mathrm{d}W_t^2, \mathrm{d}N_t \rangle = 0$.

If one allows for jumps in the model, the short end of the volatility surface will be better captured by the model while random jumps have no real effect in the long run. Jumps can increase stock price kurtosis for short maturities to reasonable levels (Bakshi et al. (1997) or Das and Foresi, 1996). For longer dated options the jump effect disappears and the fits of the Heston model and the Bates model are comparable. To assure that the volatility process always stays positive, the Feller condition should be fulfilled. The underlying price process will stay positive even if jumps occur, because the jumps are modeled as multiplicative jumps regarding the logarithm of the underlying price process. Therefore, the worst case scenario is a jump to zero which results in an instant default of the modeled underlying.

European plain-vanilla options can be priced analytically by the same technique as for the Heston model, with only minor adaptions. Hence, a model calibration to plain-vanilla prices is possible in a similar way.

2.5 Model Risk

Model risk is a well known problem in academia and practice and even though it exists a broad strand of literature about model risk in finance there is no clear cut definition.[1] In the case of derivative pricing, model risk can be described as the risk that different models yield different prices for the same exotic product even though the models are calibrated to observable prices of plain-vanilla products and reproduce these prices perfectly. This general definition of model risk can be further decomposed since model risk is manifold. Tunaru (2015) presents a comprehensive overview of model risk regarding the pricing of derivatives, defining five sub-categories of model risk. The first one is parameter estimation risk which results from wrong or inaccurate parameter estimations. The next component is model selection risk which deals with the risk resulting from not choosing the best model if the correct model class is known. The most common definition of model risk in general is model identification risk, which is the risk stemming from the usage of an suboptimal model and without knowledge about the true model class. The fourth category is model implementation risk and the last one is model protocol risk which has its origin in different understandings of the model or its parameters by two or more people. The last two categories can be seen as a kind of operational risk and are therefore excluded from the following model risk definition. The first three categories are kept and the parameter estimation risk is further divided into two sub-categories, calibration risk and recalibration risk. Calibration risk describes the risk resulting from different calibration methods or the usage of different calibration instruments. Meanwhile, recalibration risk deals with a time component and captures the risk originated in periodical recalibration using the same methods and comparable calibration instruments. An overview of model risk and its components is given in figure 1.

[INSERT FIGURE 1 ABOUT HERE]

This study focuses on model identification risk and parameter estimation risk. It is not possible to analyze model selection risk since no optimal class of models for the purpose of derivative pricing is known. To measure model risk, the approach proposed by Cont

[1] See for example Derman (1996), Green and Figlewski (1999), Kerkhof et al. (2010), or Branger and Schlag (2004) for the examination of model risk in finance.

(2006) is followed, who introduced a coherent model risk measure. Assume a set of models $Q = \{Q\}_{i=1}^{n}$ and a set of calibration instruments $H = \{H\}_{i=1}^{m}$ with corresponding bid and ask prices C_i^{bid} and $C_i^{ask}, i = 1, \ldots, m$. The models are calibrated in a way that all resulting model prices are between the bid and ask prices. Then, a coherent measure of model risk for a contingent claim C is given by

$$\psi_Q(C) = \overline{\pi}(C) - \underline{\pi}(C), \tag{10}$$

with $\overline{\pi} = \sup_{i \in \{1,\ldots,n\}} \mathrm{E}_{Q_i}(C)$ and $\underline{\pi} = \inf_{i \in \{1,\ldots,n\}} \mathrm{E}_{Q_i}(C)$ which is nothing else than the lowest and the highest theoretical model price. Therefore, it is sufficient to take the range of model prices to get a measure of model risk. Obviously, model risk measured in empirical studies is just a lower bound of the real model risk since not every model can be used. Due to this restriction, one has always to consider the set of employed models in model risk analyzes.

3 Model Calibration

3.1 Monte-Carlo Simulation

Regarding the models described in the preceding sections, it is not possible to use closed form solutions to price barrier options. Therefore, it is necessary to use numerical methods for option pricing, in particular, a Monte-Carlo approach. To simulate the underlying price process, an Euler discretization is applied for all models. Regarding the Heston model and the Bates model, the volatility process is discretized using a Milstein scheme.[2] To improve the convergence order of the used simulation approach, we use the Richardson extrapolation method, which yields a discretization with a convergence order of two without applying higher order discretization schemes.[3] The Richardson extrapolation uses two different outputs of the Monte-Carlo simulation, one with step size n and one with step size $\frac{n}{2}$ and combines them to reduce the included error. The improved estimator is $f_\infty = 2f_n - f_{\frac{n}{2}}$, where f_n and $f_{\frac{n}{2}}$ are the estimates using a step size of n and $\frac{n}{2}$, respectively. It is possible that the underlying price process crosses the barrier between two consecutive simulation points, leading to upward biased results. To minimize this problem, the

[2]For details on the discretization methods see for example Gatheral (2006).

[3]For details see for example Glasserman (2004).

probability of a barrier crossing between two succeeding steps i and $i+1$ is calculated. This probability can be derived as

$$\pi_i = \begin{cases} 1 & \text{if } \min\{S_i; S_{i+1}\} < H, \\ exp\left(-2\frac{(S_i-H)(S_{i+1}-H)}{S_i^2\sigma_i^2\Delta t}\right) & \text{else,} \end{cases} \tag{11}$$

where S_i and S_{i+1} are the underlying prices at step i and $i+1$, respectively. The barrier level is denoted by H, the volatility for the interval between step i and $i+1$ is called σ_i and Δt is the step size. The probability that the barrier is not crossed over the whole path is therefore

$$1 - \pi = \prod_i (1 - \pi_i). \tag{12}$$

This result can be used to calculate a more reliable estimator of the option price by multiplying the payoff under no barrier crossing with the calculated crossing probability. The inclusion of jumps in the Bates model requires an extension of the Monte-Carlo approach. The number of jumps and the jump size have to be simulated. Instead of simulating jump times, we simulate the number of jumps and the according jump sizes in a certain interval, although this approach can cause problems similar to the barrier crossing problem described above. Nevertheless, this second approach is computationally more efficient and the potential problems occur rarely with our calibrated parameter set. After carrying out some numerical tests it is sufficient to use 100,000 simulation paths and step sizes of $\frac{1}{100}$ years and $\frac{1}{50}$ years, respectively. This parameter choice leads to stable option price estimates and reduces the remaining simulation error to a sufficient extent.

3.2 Calibration of the Local Volatility Model

According to Dupire (1994), it exists a unique connection between European call option prices and the corresponding local volatility parameters. This relation can be expressed as

$$\sigma^2(S_t, t; S_0) = 2\left.\frac{\frac{\partial C}{\partial T} + rK\frac{\partial C}{\partial K}}{K^2\frac{\partial^2 C}{\partial K^2}}\right|_{K=S_t;T=t}, \tag{13}$$

where S_t is the price of the underlying at time t, C denotes the price of an European call option with strike K and maturity T, and r is the continuous risk-free interest rate. To apply this formula, the partial derivatives of the call option price with respect to the strike price K and the time T have to be calculated. This has to be done numerically. Due to market imperfections and noise in the market data, these derivatives tend to be unstable, leading to potentially unrealistic local volatility parameters such as a negative volatility for example. To avoid this problem, a smoothed implied volatility surface instead of option prices is used to calculate the local volatility parameters. The corresponding formula to calculate the local volatility based on the implied volatility is given by

$$\sigma^2(t, S_t) = \frac{\frac{\partial w(k,T)}{\partial T}}{1 - \frac{k}{w(k,T)}\frac{\partial w(k,T)}{\partial k} + \frac{1}{4}\left(-\frac{1}{4} - \frac{1}{w(k,T)} + \frac{k^2}{w^2(k,T)}\right)\left(\frac{\partial w(k,T)}{\partial k}\right)^2 + \frac{1}{2}\frac{\partial^2 w(k,T)}{\partial k^2}} \tag{14}$$

where $\sigma^2_{BS}(k,T)$ denotes the implied Black-Scholes variance, $w(k,T) = \sigma^2_{BS}(k,T) \cdot T$ is the total implied variance, k is the log-forward moneyness defined as $k = \log\left(\frac{K}{F_T}\right)$ with the forward price $F_T = S_t e^{r\cdot(T-t)}$.[4] To handle the problem of intractable partial derivatives, a parametric volatility surface is fitted to the given implied volatility to receive a smooth volatility surface free of static arbitrage with stable partial derivatives.[5] The iterative approach proposed by Gatheral and Jacquier (2014) is used to fit arbitrage-free volatility surfaces to the given Black-Scholes implied volatility. First, a whole arbitrage-free volatility surface is fitted to the given data, before this volatility surface is differentiated for every observable maturity to achieve a better overall fit. The starting surface is parameterized as

$$w(k, \theta_t) = \frac{\theta_t}{2}\{1 + \rho\varphi(\theta_t)k + \sqrt{(\varphi(\theta_t)k + \rho)^2 + (1-\rho^2)}\} \tag{15}$$

with $\varphi(\theta_t) = \frac{\eta}{\sqrt{\theta_t}}$ and $w(k, \theta_t)$ is a parameterized version of the total variance. While θ_t varies for different time horizons, the other two parameters ρ and η remain constant for every maturity. This surface is fitted to the data by minimizing the sum of squared distances between given implied variances and theoretical variances. The minimization is carried out with a quasi-Newton method that allows for parameter constraints introduced by Byrd et al. (1995). The resulting parametric volatility surface is arbitrage-free, but normally it does not fit the given data in a sufficient way. Therefore, some additional optimization steps are applied to achieve a better overall fit while retaining the no-arbitrage property. Since the best possible fit for every observable time horizon should be calculated,

[4] See for example Gatheral (2006) for a detailed derivation of this formula.

[5] See for example Gatheral and Jacquier (2014) for a definition of arbitrage-free volatility surfaces.

it is not useful to have common parameters for different time horizons, anymore. Therefore, the following parametrization is used for single time horizons

$$w(k, \Theta) = a + b\{\rho(k - m) + \sqrt{(k - m)^2 + \sigma^2}\} \tag{16}$$

with the parameter vector $\Theta = (a, b, \rho, m, \sigma)$. The shape of the fitted volatility smile for a given time horizon t is defined by these five parameters. By an iterative procedure which obtains the no-arbitrage property, a better overall fit than the initial parametrization is obtained. The local volatility parameters are calculated by applying Formula (14) to this volatility surface with numerical differentiation. For this purpose, in the time dimension an inter- or extrapolation is required. Following Gatheral and Jacquier (2014), the Stineman (1980) interpolation method is applied, which yields monotonically increasing values in contrast to the often used cubic spline interpolation.

3.3 Calibration of the Heston Model and the Bates Model

The calibration of the Heston model and the Bates model is carried out in a similar way. The calibration is based on the fact that there are closed-form solutions for pricing European call options in both models. Therefore, it is possible to directly minimize the differences between observable and theoretical call option prices influenced by the model parameters. The chosen objective function to be minimized is the root mean squared error between observed and theoretical call option prices

$$\sqrt{\frac{1}{n}\sum_{i=1}^{n}\left(C^i_{obs} - C^i_{model}\right)^2} \rightarrow \min \tag{17}$$

where n is the number of observed options, C^i_{obs} is the observed price of the i-th call option and C^i_{model} is the theoretical value of the i-th call option. The main problem regarding the calibration process is the existence of different local minima. Therefore, it is necessary to either use a local optimization algorithm with a reasonable initial value close to the global minimum or a global optimization algorithm. In this paper, a global optimization algorithm is applied to make sure that the global minimum is detected without the risk of sticking at a local minimum. If the parameters that correspond to a local minimum are used to price the certificates, an additional model risk, namely calibration risk, is present in the theoretical model prices.

We apply the differential evolution algorithm proposed by Storn and Price (1997) to calibrate the models, which generally leads to the global minimum.[6] The optimization has to be carried out under some constraints. First, initial limits for the parameters are used which can be seen as a minimal requirement. Second, it is assured that the boundaries are wide enough to allow for an increase or decrease of 50% for every parameter value from one day to another as long as the theoretical limits are not violated. In this way, the initial boundaries can be expanded through the second condition.[7]

If both the Heston and the Bates model are calibrated to the same data set under the same parameter constraints with the same objective function, the optimal solution regarding the Bates model should yield an objective function value smaller or equal to the optimal objective function value in the Heston case.[8] This property is useful for the validation of the calibration results. The differential evolution algorithm needs a terminal condition to stop searching. Some numerical tests lead to the finding that a number of 200 iteration steps is sufficient for the algorithm to find the optimal solution in the Heston case, no matter the market environment. Regarding the Bates model, a number of 200 iteration steps seems to be sufficient, too. Nevertheless, in nearly 10% of the trading days, the calibration method results in higher objective function values compared to the calibration of the Heston model. Therefore, the optimization process for these days is repeatedly continued with an additional 100 iteration steps until the objective function value regarding the Bates model is smaller or equal than the objective function value in the case of the Heston model.

An overview of the optimal model parameters is given in Table 1. Regarding the Heston model calibration, two different data sets are used. First, all available options are used, and second, the model is calibrated to the standard DAX options only, excluding the weekly options with short maturities. Due to the fact that the Bates model yields

[6]For the calibration of the Heston model via the differential evolution algorithm see for example Gilli and Schumann (2010) or Vollrath and Wendland (2009).

[7]The initial boundaries are given by 0 and 0.5 for v_0, θ and v_J, 0 and 5 for κ, 0 and 2 for ξ, -1 and 0.3 for ρ, 0.1 and 1.5 for λ and -0.25 and 0.1 for μ_J. There were in fact several days with an expansion of the initial boundaries, in most cases for the mean reversion speed κ.

[8]This only holds true if one allows for the parameter λ to become zero or for the parameters μ_J and v_J to become zero simultaneously, which is possible in this study.

better calibration results for short maturities, the weekly options are not excluded for the Bates model calibration. The Heston model has no such advantages for short maturities and therefore the two different data sets are considered to analyze if the fit for longer maturities improves with lesser short dated options. Additionally, it is possible to analyze the calibration risk in the case of the Heston model.

[INSERT TABLE 1 ABOUT HERE]

Table 1 reveals that no real difference is visible for the parameters v_0 and θ in both Heston model calibrations. While ξ tends to be slightly lower if the model is calibrated to fewer options, ρ tends to be a little higher. The parameter κ experiences the biggest shift. Using fewer options to calibrate the model results in a considerably lower κ regarding the mean value.

Comparing the Bates model with the Heston model results for the whole data set, it is obvious that every parameter except κ is shifted downwards, while κ itself is shifted upwards significantly. The calibration results seem to be reasonable from an economic point of view.

4 Data

4.1 Market Data

The certificates' prices are calculated using the models described in section 2. Regardless which model is applied, one needs the same input data, including the continuous risk-free interest rate, option prices of plain-vanilla options, the price of the underlying, and issuer individual credit spreads.

The continuous risk-free interest rate is derived from German government bonds. The whole yield curve is modeled with the approach proposed by Svensson (1994). The model parameters are provided by the Deutsche Bundesbank on a daily basis via its internet site. Daily settlement prices for European put and call options written on the DAX and traded at the EUREX are used for the model calibrations. Standard DAX options with expiry at the maturity month's third Friday as well as weekly DAX options are used. The weekly option data is provided by the Karlsruher Kapitalmarktdatenbank, while the basic option data is provided by Thomson Reuters EIKON.

The underlying prices of the DAX performance index are observed simultaneously to the quotes of the certificates to guarantee synchronicity. DAX tick data is provided by the Karlsruher Kapitalmarktdatenbank as well.

Daily credit spreads for all issuers are obtained via Thomson Reuters EIKON.

Models are calibrated to settlement prices of EUREX options, which are settled at 5:30 p.m. on every trading day. Therefore, the last available certificate bid and ask prices before 5:30 p.m. are obtained to guarantee the best possible synchronicity. The bid and ask prices from EUWAX are provided by VWD.

4.2 Certificates Data

The data set consists of classic and capped bonus certificates written on the DAX, issued between May 2, 2014 and December 30, 2015. Only certificates which mature before December 30, 2015 are included to make sure that an analysis over the whole life cycle is possible. Issuers with a very small amount of outstanding products are excluded. All information about the certificates such as barrier and bonus level, issue date, maturity date and issuer are provided by Deriva GmbH.

The product selection results in 40,670 certificates of 11 issuers. The majority of these products are capped bonus certificates with a number of 26,373 in contrast to 14,297 classic bonus certificates. A detailed overview of the data set is given in Table 2.

[INSERT TABLE 2 ABOUT HERE]

The certificates have an average time to maturity of 168 days with a minimum of 7 days and a maximum of 590 days. The average moneyness at issuance, which is the relative distance to the bonus level, is given by 45.94% for classic bonus certificates and 61.85% for capped bonus certificates. The average cushion at issuance is equal to 15.31% for classic bonus certificates and 14.87% for capped bonus certificates. The cap and the bonus level are equal for all capped bonus certificates.

The number of certificates per issuer is unevenly distributed. While the data set for Citigroup consists of about 500 certificates, the data set of Goldman Sachs contains more than 10,500 certificates. Unicredit, HSBC and Vontobel have issued only capped bonus certificates during the period under study.[9] The time to maturity is around 0.5 years for most issuers and product types. Longer time horizons are visible for DZ Bank and BNP regarding classic bonus certificates and Citigroup and BNP regarding capped bonus certificates. Shorter time spans are observable for Goldman Sachs regarding classic bonus certificates and Deutsche Bank, Goldman Sachs, Société Générale and HSBC for capped bonus certificates.

Regarding the moneyness, it is obvious that Citigroup uses by far the lowest bonus levels for both types of certificates. Most issuers show comparable values for both types with Deutsche Bank and Société Générale being the exceptions. Société Générale uses significantly lower bonus levels while Deutsche Bank tends to use extraordinary high values. Only HSBC issues capped bonus certificates with even higher bonus levels than Deutsche Bank. These two issuers have a lot of certificates in their respective portfolio with bonus levels between 40,000 and 50,000 index points.

The huge range of the bonus levels reveals different understandings or marketing approaches regarding bonus certificates. While a higher bonus level suggests that the retail investor gets a defined amount of money as a bonus if the barrier is still in place at maturity, the very low bonus level used by Citigroup suggests a capital protection with the chance to earn an additional bonus. Even though the products are constructed in identical ways, they can be sold under different understandings which is an example of the framing effect.

The cushion does not vary much between the issuers and there are no substantial deviations between the two types of certificates regarding the single issuers. The lowest cushion and thereby the highest barrier is used by Goldman Sachs with nearly 10%. The result of the combination of extremely high bonus levels and really low cushions are products with very high risk.

[9]HSBC issued 32 classic bonus certificates which are excluded from the data set since the sample is too small.

5 Model Risk Analysis

5.1 Model Identification Risk

In this section the extent of model identification risk is analyzed. Since there is no real parameter estimation risk in the case of the Heston model, both Heston prices are incorporated in the analysis. Model identification risk ψ of a given certificate C is measured by the range of theoretical model prices in relation to the mean of the product's bid and ask prices:

$$\psi(C) = \frac{\max\{C_{Hfull}, C_{Hpartial}, C_B, C_L\} - \min\{C_{Hfull}, C_{Hpartial}, C_B, C_L\}}{0.5 \cdot (C_{bid} + C_{ask})} \tag{18}$$

with $C_{Hfull}, C_{Hpartial}, C_B, C_L, C_{bid}$ and C_{ask} the price of the certificate in the Heston model calibrated to the full data set, the Heston model calibrated to the partial data set, the Bates model, the local volatility model, the bid price, and the ask price, respectively. A histogram of the resulting model risk for classic and capped bonus certificates combined is given in Figure 2.[10]

[INSERT FIGURE 2 ABOUT HERE]

Obviously, a lot of situations with low model risk exist, but there are situations with a substantial amount of model risk as well. It is obvious that model risk can reach a significant extent and that it is necessary to account for it in pricing purposes and investment decisions. The mean model risk is 1.11%. More key figures of the distribution are given in Table 3.

[INSERT TABLE 3 ABOUT HERE]

Regarding the construction of bonus certificates, model risk might depend on product parameters. The main features are (i) the relative distance to the barrier which, is called cushion:

$$Cushion_t = \frac{S_t - B}{S_t}, \tag{19}$$

[10] All results presented in this section were derived for classic and capped bonus certificates separately. Since the differences in the results between both types of bonus certificates were negligible, only aggregated results are reported.

with S_t the underlying price at time t and the product's barrier B, and (ii) the remaining time to maturity. Since the figures should be comparable for all bonus certificates, the relative product lifetime is defined as

$$Lifetime_t = \frac{t - t_0}{T - t_0}, \tag{20}$$

where t_0 denotes the issuance and T the maturity date.

The dependency on the product's lifetime is shown in Figure 3 (a). There is no striking relation between model risk and the relative lifetime except for the last 15 percent of the products' lifetime. Model risk declines in this area and nearly disappears at maturity. This decline is expected since the payoff at maturity becomes less uncertain at the end of a product's lifetime, at least in the majority of possible situations. This fact holds especially true for capped bonus certificates close to maturity, if the underlying lies above the bonus level.

[INSERT FIGURE 3 ABOUT HERE]

Regarding the relative distance to the barrier, there is a more pronounced relation which is given in Figure 3 (b). No significant amount of model risk is visible for a higher distance to the barrier, since the knockout is unlikely regardless the model. As the underlying approaches the barrier, the model risk rises with the maximum close to the barrier at a relative distance of about two percent. Directly above the barrier, the model risk experiences a sharp decline. This pattern can be explained with the probability of reaching the barrier, which takes different values in different models and causes broader price ranges. At a point close to the barrier, these probabilities are very high for every model, leading to the sharp decline in model risk.

To analyze the common influence of the aforementioned factors on model risk, polynomial regressions including interaction terms are performed to find the best possible fit to the data set. The resulting model risk surface in dependence of the relative lifetime and the relative distance to the barrier is given in Figure 3 (c). It is obvious that model risk tends to decrease with an increasing distance to the barrier regardless the remaining lifetime. The surface is more or less flat with the exception in the area where short distances to the barrier and short remaining times to maturity coincide. In this case, model risk reaches its highest values by far. On first sight this seems to be surprising since a decrease in model risk for products close to maturity is shown above. Since this is true in general, a special case is given by products close to the barrier. Different model prices and therefore model risk result from different barrier hit probabilities. Regarding products close to the barrier, the longer the time to maturity the more probable are high barrier hit probabilities for all models under study, since a minor downside movement of the underlying is sufficient to hit the barrier. Regarding products close to the barrier and close to maturity, there are bigger differences between barrier hit probabilities, since the different dynamics of the underlying in the different models result in diverging short term behavior. The difference between model prices in this situation is mainly introduced by the distance between the local volatility model and the other two models. The local volatility model experiences far more barrier hits and therefore yields lower prices. The reason for this is given by the local volatility surface, since there is a really high level of volatility for short maturities which causes the underlying to experience higher deviations and thereby increases the risk of hitting the barrier.

Model risk arises from the range of barrier hit probabilities as was already mentioned above. Due to this fact, we additionally calculate these probabilities for all three models. These probabilities are the number of simulated paths that cross the respective barrier divided by the total number of simulated paths. The model risk in dependence of the range of barrier hit probabilities is shown in Figure 3 (d). Since all models are calibrated to the same option data they will produce comparable distributions of the underlying price at maturity. Therefore, price differences mainly occur due to different behavior along the paths. That is the explanation for the nearly absence of model risk for really small ranges of barrier hit probabilities. The model risk appears to grow linearly with an increasing range of barrier hit probabilities.

5.2 Model-by-Model Analysis

The results stated above reveal that a significant amount of model risk is measurable and that there exist situations which are prone to high model risk. This section analyzes the question which model drives the model risk.

Since the Bates model is an extension of the Heston model it can be expected that the resulting model prices are close to each other. If this assumption holds true, the results given above are mainly a consequence of the difference between the local volatility model and the two other models. To verify this assumption, the pairwise differences between the three models are analyzed. The key figures of the distribution of the pairwise differences are given in Table 3. It becomes clear that no model dominates another model for every observation. Nevertheless, it is possible to find a hierarchy between the models, since a clear relation between different model prices is visible for the majority of observations. Both the Heston model and the Bates model yield higher theoretical product prices than the local volatility model, on average. The Bates model results in the highest product prices. The mean difference between the theoretical prices is minimal for the comparison of the Heston model and the Bates model with 0.41%. The mean difference between the local volatility model and the two other models is bigger with 0.56% compared to the Heston model and 0.97% compared to the Bates model.

Regarding the influencing factors on model risk, the main difference between the local volatility model and the two stochastic volatility models is caused by price deviations for certificates close to the barrier, which can be seen in Figure 4. As already stated above, this deviation can be explained by the higher volatility level for short maturities and far away from the at-the-money point. Especially for DAX values far below the real DAX value at the time of calibration, the local volatility takes on some significantly higher values than the overall local volatility level. This is the most crucial situation, since a certificate tends to be close to the barrier in this situation and therefore a higher volatility which corresponds to higher absolute price movements results in a higher barrier-hit probability and consequently a lower product price.

[INSERT FIGURE 4 ABOUT HERE]

In the case of the two stochastic-volatility models, a less pronounced pattern is visible regarding the dependance on the distance to the barrier. There is still an increase in model risk, followed by a decline as the underlying approaches the barrier, but the maximum is substantially further away from the barrier compared to the results stated above at a relative distance of about 10%. One reason for this behavior can be the presence of jumps in the underlying price process. Since the expected return in the risk neutral world has to be equal to the continuous risk free interest rate, the drift rate in the Bates model has to be adjusted due to the presence of jumps. The average jump size is negative and therefore the negative expected return caused by jumps has to be compensated via a higher drift rate. This higher drift rate in combination with the rarely occurring jumps leads to lower barrier hit probabilities and potentially higher payoffs and consequently to higher product prices.

Regarding the relative lifetime, the results stated above for the comparison of all three models hold true for the two model case, as well. The relative lifetime has no obvious impact on the model risk except for the last few percent of the relative lifetime where model risk converges to zero.

5.3 Recalibration Risk

In this section the evolution of model risk over time is analyzed, that is, recalibration risk. To decide whether model risk is stable over time or not, the daily differences in model risk on a single product basis are determined. Let $\rho_{i,t}$ be the model risk for product i at trading day t, then the change of model risk from one trading day to the next is given by $\rho_{i,t} - \rho_{i,t-1}$. Some figures about the distribution of these daily changes of model risk are given in Table 3. There is not a lot of variation in the data with the mean change equal to -0.003% and a standard deviation of 0.49%. These results show that model risk remains stable over time and is not significantly effected by the daily recalibration of the models. Nevertheless, there seem to exist situations with a little higher deviation between consecutive observations of model risk for single products. Since these higher deviations are rare events, in general the model risk for the whole portfolio of bonus certificates does not vary significantly on a daily basis.

Regarding the different model parameters, especially in the case of the Heston model and the Bates model, the results are contrary. The model parameters vary significantly over time, which is necessary to achieve the best fit to the plain vanilla option data. Since these varying parameters have no significant influence on the extent of model risk for bonus certificates, it is not necessary to account for model risk induced by periodical recalibration as long as the same calibration procedure is used.

5.4 Calibration Risk

Calibration risk is analyzed in the case of the Heston model. Therefore, the resulting optimal parameter sets for the Heston model based on two different sets of calibration instruments are analyzed. The model is calibrated to all available options including weekly options on one hand and only to standard options on the other hand. An overview of the resulting optimal parameter sets is given in Table 1. Obviously, there are no substantial differences between the optimal model parameters. The only parameter with a slightly bigger shift is the mean-reversion speed κ, which experiences a decline of about 7.6% if only the standard options are used as calibration instruments. Since the influence of the mean-reversion speed on option prices is least pronounced for all model parameters one would not expect a significant amount of model risk induced through this difference.[11] Key figures about the distribution of the resulting model risk are given in Table 3. There is no significant extent of calibration risk in the case of the Heston model, since nearly all observations lie in an interval from -0.2% to 0.2% with a mean of -0.04% and a median of -0.01%.

6 Model Choice of the Issuers

In this section we analyze, which models are used by the issuers to calculate certificates' prices. Therefore, the issuer bid-ask prices are compared to the theoretical product prices. It is not possible to directly compare these values, since the market prices include unknown issuer margins. Additionally, the product prices calculated on consecutive days are probably highly correlated. To get rid of the two aforementioned problems, daily returns are calculated for every product with each theoretical model and additionally with the actual mid price of the issuers.

To identify the model which issuers actually use, we apply a regression approach similar to Baule and Tallau (2011). The returns based on issuer market data $r_{i,t}^{market}$ are explained by the theoretical returns $r_{i,t}^{model}$ under the incorporation of daily fixed effects α_t for every

[11] We also carried out sensitivity analyses regarding the influence of the different model parameters on the prices of bonus certificates. These numerical analyses reveal that a variation of the mean reversion speed has the smallest impact on product prices.

day t:

$$r_{i,t}^{market} = \sum_t \alpha_t + \beta \cdot r_{i,t}^{model} + \epsilon_{i,t}. \tag{21}$$

These regressions are performed separately for each issuer. If the market price is a combination of a certain model price under the incorporation of default risk and an additional issuer margin, the model's corresponding beta factor is (close to) one and the error term is (close to) zero.

Empirical literature on issuer pricing of certificates shows strong evidence that issuer margins decrease over a product's lifetime. Therefore, the daily returns based on market prices should be lower than their theoretical counterparts. This effect should result in negative alpha coefficients, on average.[12]

Based on the regression coefficients, we can infer the likelihood that issuers prefer a specific model over another. This likelihood is the larger, the closer the beta coefficient to one and the closer the error terms to zero (equivalently, the larger the regression R^2). However, even if the issuer actually uses a certain model, we cannot expect perfect results because of a number of potential sources of deviation, for example the calibration method, the calibration instruments, or the applied discretization scheme. Furthermore, statistical noise, for instance due to a non-perfect synchronization of prices, will blur the results. Nevertheless, the appropriate model should yield the lowest deviations from the ideal results. In particular, we judge by the standard deviation of the residuals to decide whether a given model is close to that used by the issuer.

The regression results, separately for classic bonus certificates and capped bonus certificates, are given in Table 4.

[INSERT TABLE 4 ABOUT HERE]

Regarding classic bonus certificates, except for Goldman Sachs, the adjusted R^2 is above 96% for at least some models for all issuers. This result is confirmed by the standard deviation of the residuals which are fairly low for all issuers except for Goldman Sachs. These observations are a first indication that the issuers indeed use a model close to the theoretical models of our study.

[12]Since the alpha coefficients absorb all other effects not explained by the simple linear model, we cannot expect all coefficients to be negative.

The identification of the most likely model is mainly based on the standard deviations. The standard deviations for the Heston model reach the lowest value for every issuer. The only other model with fairly low values for all issuers is the Bates model. The Bates model has only a slight advantage over the Heston model regarding Goldman Sachs, which was already identified as an issuer who potentially uses a different model than the ones under study. Based on the standard deviations, the local volatility model seems to be no potential choice for any issuer.

Based on the beta coefficients, again the only strong candidate to be used by the issuers is the Heston model. The Bates model yields fairly good results as well but the Heston model is always at least slightly closer to one. In combination with the observations regarding the standard deviations, we conclude that the Heston model is more likely to be used than the Bates model. The beta coefficients for the local volatility model are significantly lower than one, so this model is no appropriate candidate to be the one used by the issuers.

In summary, none of the analyzed models seems to be used by Goldman Sachs. The remaining issuers most likely use the Heston model or at least a Heston-like model, that means, a model with similar properties as in the Heston case. This could be an extension of the Heston model or a different modeling approach which results in comparable properties. The Bates model is a probable choice, too, but the results are always slightly inferior to the Heston results. The local volatility model is pretty unlikely to be used by the issuers.

The regression results for capped bonus certificates show mainly comparable results. Regarding the adjusted R^2, Goldman Sachs values are lower than the majority of the other issuers corresponding values. In contrast to the results for classic bonus certificates, the adjusted R^2 for UBS is extremely low for all models with values lower than 70%. This indicates that UBS uses different models or extremely different pricing policies for the two types of bonus certificates. Since the product features are comparable for classic and capped bonus certificates issued by UBS, the difference does not seem to be based on the product construction.

Regarding the standard deviation of the residuals, the observations for classic bonus certificates can be confirmed. The Heston model yields the lowest values with the Bates model close behind. The local volatility model seems not to be used by the issuers.

The results based on the analysis of the beta coefficients can be confirmed, too. The Heston model yields the best values and for the majority of the issuers they are close to one. An exception is HSBC which has a maximal beta coefficient of about 70%. This could be seen as an indication that HSBC uses a different model or a different pricing policy than the assumed one. This impression is confirmed by the standard deviations of the residuals which are among the highest values behind only Goldman Sachs and UBS.
In summary, Goldman Sachs, UBS, and maybe also HSBC use other models than the ones under study. All remaining issuers seem to use the Heston model. The Bates model is always slightly worse and the local volatility model is unlikely to be used by the issuers for pricing purposes.
As robustness checks, the calculations were repeated under the incorporation of default risk and without the use of daily fixed effects. The results remain nearly the same and therefore they are not reported.
So we can answer the questions raised in the introduction as follows:

- Issuers prefer models with stochastic volatility over deterministic volatility.
- Issuers prefer pure-diffusion models over models with jumps.
- Most issuers agree in their model preferences. Some issuers however apply completely different pricing policies.

7 Conclusion

In this paper, the extent of model risk present in tradable bonus certificates written on the DAX is analyzed. Therefore, theoretical model prices are calculated and the resulting range of prices is a measure of model risk. This measure is mainly influenced by the considered set of models. Since it is impossible to use every available model, the analyzed model risk has to be seen as a lower bound of the real model risk. Often, the starting point for model selection is the famous Black-Scholes model. Although the stock price dynamics with constant volatility are not overly realistic, the model yields reasonable product prices if the volatility is calibrated to similar traded derivatives. Since barrier options are not traded liquidly, the Black-Scholes model is not applied in this study. The chosen models are the local volatility model by Dupire, the stochastic volatility model by Heston, and the jump-diffusion model with stochastic volatility by Bates. This set of models is frequently used in academia and practice and therefore the resulting model risk seems to be an appropriate measure for practical purposes. Furthermore, this seems to be a natural choice, since the Black-Scholes model as a starting point is consequently expanded by the more sophisticated models to incorporate empirical findings about stock price and volatility behavior.

More than 40,000 certificates from 11 issuers with more than 3 million single observations are analyzed. The result is an amount of model risk of about 1.1% on average. This figure is substantial, compared to average issuer margins in the range of 2%–3% (Baule and Tallau, 2011). The extent of model risk depends on different factors. The highest model risk is present if a certificate is close to maturity and close to the barrier. Regarding the barrier, model risk tends to rise if the underlying approaches the barrier and reaches it maximum close to the barrier at a relative distance to the barrier of about 2%. If the underlying further approaches the barrier, model risk experiences a sharp decline, since the probability of a potential barrier hit converges for the models under study and the model risk is directly linked to the uncertainty about barrier hits.

A second analysis relates to the question which models are actually used by the issuers. To analyze this question, the returns based on market prices are regressed on the corresponding returns based on different model prices. This analysis reveals that the issuers tend to use models related to the Heston model. In some cases the Bates model is a possible choice, too. Regarding the local volatility model, there is no evidence that it is used by the issuers for the purpose of the pricing of bonus certificates. Only a few issuers seem to use completely different models or different pricing policies.

References

An, Y. and W. Suo (2009). An empirical comparison of option-pricing models in hedging exotic options. *Financial Management 38*, 889–914.

Andersen, T. G., L. Benzoni, and J. Lund (2002). An empirical investigation of continous-time equity return models. *Journal of Finance 57*, 1239–1284.

Bakshi, G., C. Cao, and Z. Chen (1997). Empirical performance of alternative option pricing models. *Journal of Finance 52*, 2003–2049.

Bakshi, G. and D. Madan (2000). Spanning and derivative-security valuation. *Journal of Financial Economics 55*, 205–238.

Barndorff-Nielsen, O. E. and N. Shephard (2001). Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 63*, 167–241.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Review of Financial Studies 9*, 69–107.

Baule, R. (2011). The order flow of discount certificates and issuer pricing behavior. *Journal of Banking & Finance 35*, 3120–3133.

Baule, R., O. Entrop, and M. Wilkens (2008). Credit risk and bank margins in structured financial products: Evidence from the German secondary market for discount certificates. *Journal of Futures Markets 28*, 376–397.

Baule, R. and C. Tallau (2011). The pricing of path-dependent structured financial retail products: The case of bonus certificates. *Journal of Derivatives 18*, 54–71.

Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy 81*, 637–654.

Branger, N. and C. Schlag (2004). Model risk: A conceptual framework for risk measurement and hedging. Working Paper, Goethe University Frankfurt.

Breeden, D. T. and R. H. Litzenberger (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business 51*, 621–651.

Byrd, R. H., P. Lu, J. Nocedal, and C. Zhu (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing 16*, 1190–1208.

Cont, R. (2006). Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance 16*, 519–547.

Cox, J. C. and S. A. Ross (1976). The valuation of options for alternative stochastic processes. *Journal of Financial Economics 3*, 145–166.

Das, S. R. and S. Foresi (1996). Exact solutions for bond and option prices with systematic jump risk. *Review of Derivatives Research 1*, 7–24.

Derman, E. (1996). Model risk. *Goldman Sachs quantitative strategies research notes* (April), 1–11.

Derman, E. and I. Kani (1994). Riding on a smile. *Risk 7*(2), 32–39.

Dumas, B., J. Fleming, and R. E. Whaley (1998). Implied volatility functions: Empirical tests. *Journal of Finance 53*, 2059–2106.

Dupire, B. (1994). Pricing with a smile. *Risk Magazine 7*, 1–10.

Engelmann, B., M. R. Fengler, and P. Schwendner (2009). Hedging under alternative stickiness assumptions: An empirical analysis for barrier options. *Journal of Risk 12*(1), 53–77.

Entrop, O., G. Fischer, M. McKenzie, M. Wilkens, and C. Winkler (2016). How does pricing affect investors' product choice? Evidence from the market for discount certificates. *Journal of Banking & Finance 68*, 195–215.

Eraker, B. (2004). Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance 59*, 1367–1403.

Feller, W. (1951). Two singular diffusion problems. *Annals of Mathematics 54*, 173–182.

Gatheral, J. (2006). *The volatility surface: A practitioner's guide.* New Jersey: Wiley.

Gatheral, J. and A. Jacquier (2014). Arbitrage-free SVI volatility surfaces. *Quantitative Finance 14*, 59–71.

Gilli, M. and E. Schumann (2010). Calibrating the Heston model with differential evolution. In C. Di Chio, A. Brabazon, G. A. Di Caro, M. Ebner, M. Farooq, A. Fink, J. Grahl, G. Greenfield, P. Machado, M. O'Neill, E. Tarantino, and N. Urquhart (Eds.), *Applications of Evolutionary Computation*, pp. 242–250. Berlin: Springer.

Glasserman, P. (2004). *Monte Carlo methods in financial engineering*. New York: Springer.

Green, T. C. and S. Figlewski (1999). Market risk and model risk for a financial institution writing options. *Journal of Finance 54*, 1465–1499.

Guillaume, F. and W. Schoutens (2012). Calibration risk: Illustrating the impact of calibration risk under the Heston model. *Review of Derivatives Research 15*, 57–79.

Hagan, P. S., D. Kumar, A. S. Lesniewski, and D. E. Woodward (2002). Managing smile risk. *Wilmott Magazine 1*, 84–108.

Henderson, B. J. and N. D. Pearson (2011). The dark side of financial innovation: A case study of the pricing of a retail financial product. *Journal of Financial Economics 100*, 227–247.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies 6*, 327–343.

Hirsa, A., G. Courtadon, and D. B. Madan (2003). The effect of model risk on the valuation of barrier options. *Journal of Risk Finance 4*, 47–55.

Hull, J. and W. Suo (2002). A methodology for assessing model risk and its application to the implied volatility function model. *Journal of Financial and Quantitative Analysis 37*, 297–318.

Hull, J. and A. White (1995). The impact of default risk on the prices of options and other derivative securities. *Journal of Banking & Finance 19*, 299–322.

Kerkhof, J., B. Melenberg, and H. Schumacher (2010). Model risk and capital reserves. *Journal of Banking & Finance 34*, 267–279.

Madan, D. B., P. P. Carr, and E. C. Chang (1998). The variance gamma process and option pricing. *European Finance Review 2*, 79–105.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics 3*, 125–144.

Nalholm, M. and R. Poulsen (2006). Static hedging and model risk for barrier options. *Journal of Futures Markets 26*, 449–463.

Schertler, A. (2016). Pricing effects when competitors arrive: The case of discount certificates in Germany. *Journal of Banking & Finance 68*, 84–99.

Schoutens, W., E. Simons, and J. Tistaert (2004). A perfect calibration! Now what? *Wilmott Magazine*, 66–78.

Stineman, R. W. (1980). A consistently well-behaved method of interpolation. *Creative Computing 6*(7), 54–57.

Stoimenov, P. A. and S. Wilkens (2005). Are structured products "fairly" priced? An analysis of the German market for equity-linked instruments. *Journal of Banking & Finance 29*, 2971–2993.

Storn, R. and K. Price (1997). Differential evolution – A simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization 11*, 341–359.

Svensson, L. E. O. (1994). Estimating and interpreting forward interest rates: Sweden 1992-1994. Working Paper, Stockholm University.

Tunaru, R. (2015). *Model risk in financial markets: From financial engineering to risk management*. Singapore: World Scientific Publishing Company.

Vollrath, I. E. and J. Wendland (2009). Calibration of interest rate and option models using differential evolution. Working Paper, FINCAD Corporation.

Wallmeier, M. and M. Diethelm (2009). Market pricing of exotic structured products: The case of multi-asset barrier reverse convertibles in Switzerland. *Journal of Derivatives 17*(2), 59–72.

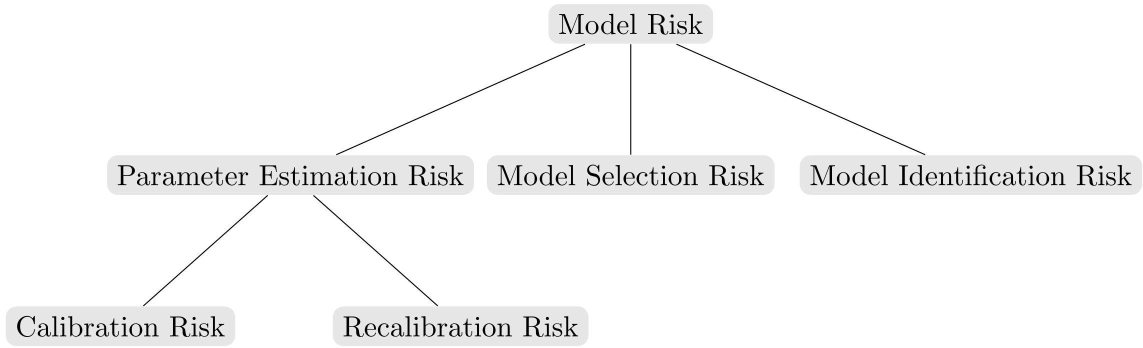

Figure 1. *Overview of model risk and its components.*

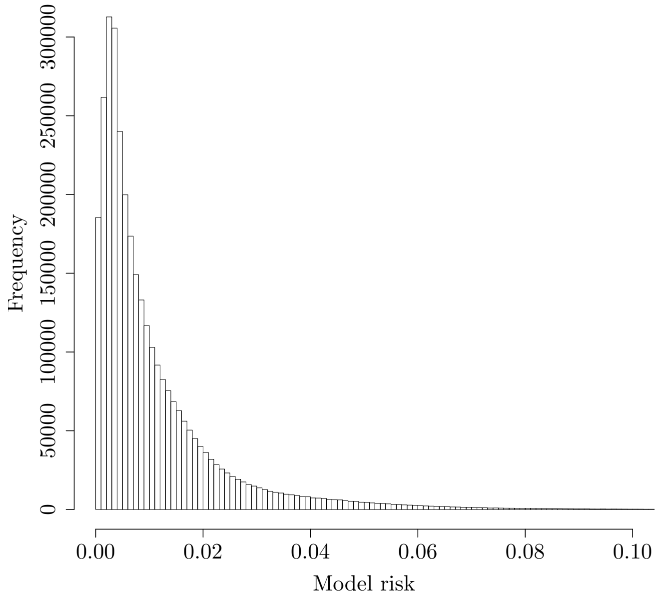

Figure 2. *Distribution of empirical model risk in percent. Model risk is defined as the difference between maximum and minimum model prices of the local volatility model, the Heston model (with two different calibrations), and the Bates model. The data base covers 40,670 classic and capped bonus certificates from 11 issuers observed between May 2014 and December 2015.*

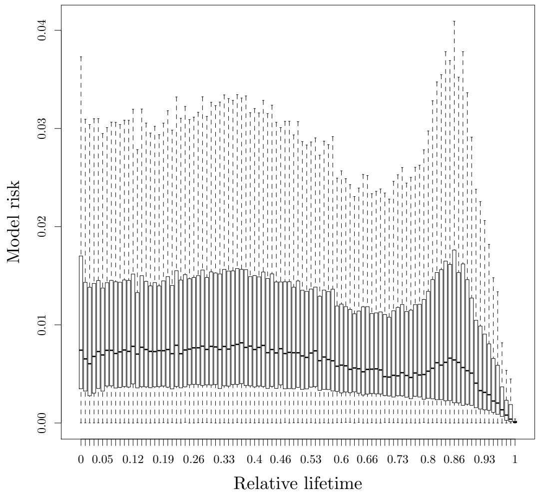

(a)

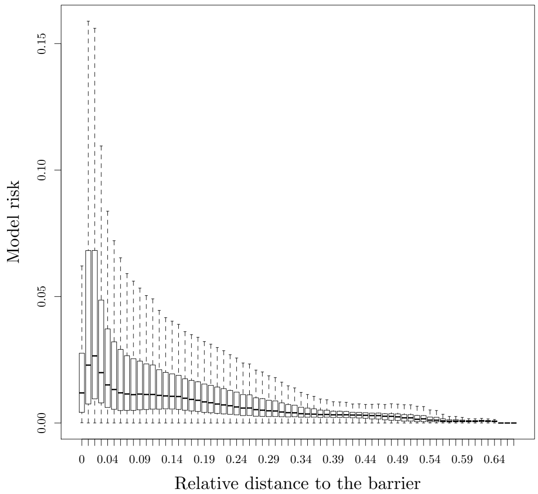

(b)

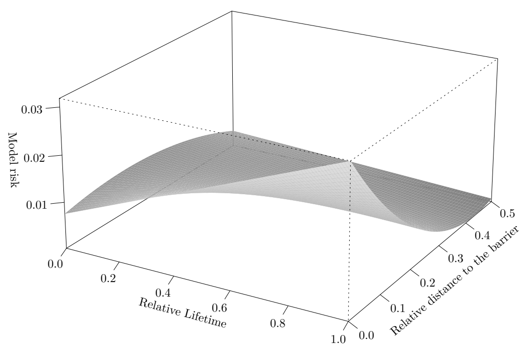

(c)

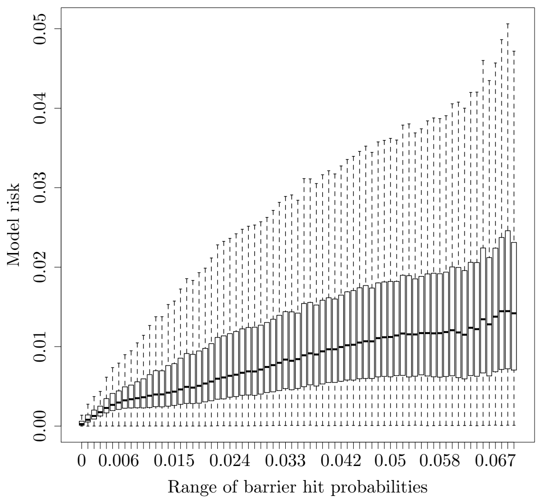

(d)

Figure 3. *Model risk in dependence of different influencing factors. Subfigures (a) and (b) refer to relative lifetime and relative distance to the barrier, respectively, Subfigure (c) to both factors. Subfigure (d) shows model risk in dependence of the range of barrier hit probabilities. Subfigures (a), (b), and (d) show a range of boxplots, with the upper whiskers limited by 1.5 times the inter-quartile range.*

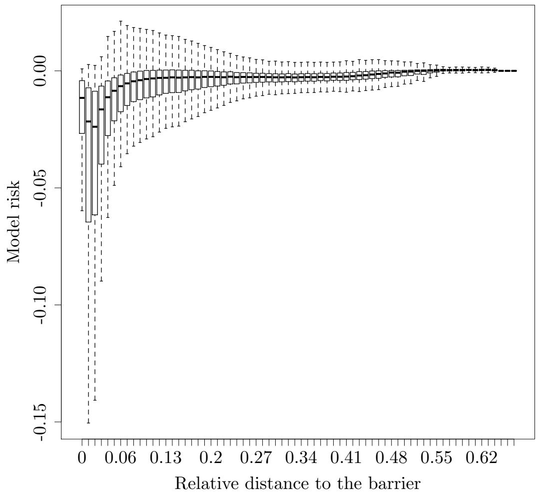

(a) *local volatility minus Heston*

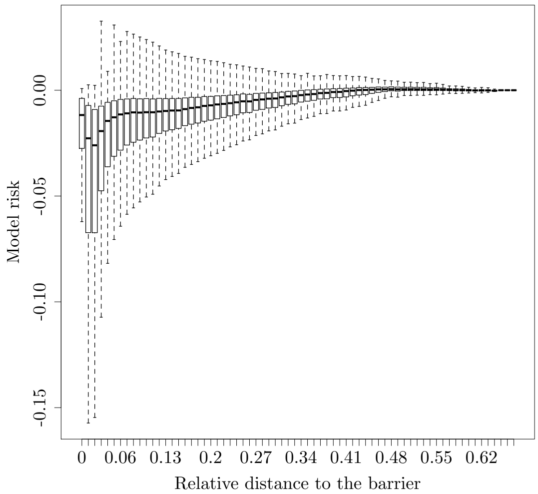

(b) *local volatility minus Bates*

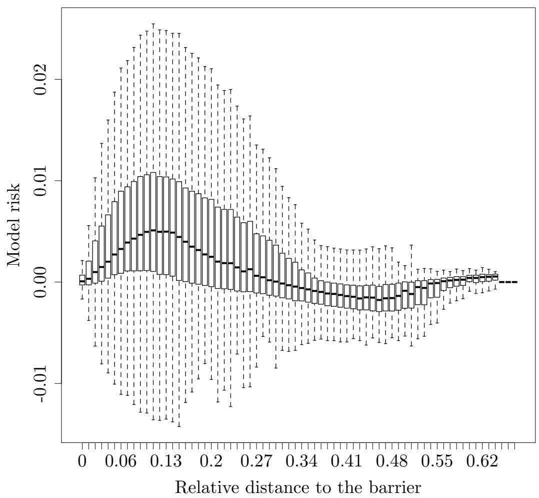

(c) *Bates minus Heston*

Figure 4. *Pairwise differences between the different models in dependence of the relative distance to the barrier. Each subfigure shows a range of boxplots, with the upper whiskers limited by 1.5 times the inter-quartile range.*

	Heston full			Heston partial			Bates		
Parameter	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max
v_0	0.043	0.011	0.171	0.042	0.012	0.168	0.041	0.009	0.176
θ	0.047	0.036	0.058	0.047	0.036	0.058	0.032	0.019	0.056
κ	2.722	0.966	11.230	2.514	0.951	9.300	3.998	0.353	31.880
ξ	0.482	0.333	0.979	0.469	0.325	0.889	0.350	0.135	1.005
ρ	-0.771	-0.997	-0.556	-0.762	-0.949	-0.543	-0.865	-0.999	-0.559
λ	-	-	-	-	-	-	0.167	0.100	1.316
μ_J	-	-	-	-	-	-	-0.125	-0.249	0.098
v_J	-	-	-	-	-	-	0.280	0.000	0.479

Table 1. *Results of model calibrations. For the Bates model, the Heston model "full" (calibrated to every available option data), and the Heston "partial" model (with the exclusion of short dated weekly options), average, minimum, and maximum optimal parameter values over the period of 422 trading days are reported.*

	Classic Bonus Certificates				Capped Bonus Certificates			
Issuer	# Certificates	Maturity	Moneyness	Cushion	# Certificates	Maturity	Moneyness	Cushion
Citigroup	322	0.57	0.34%	24.85%	209	0.69	2.97%	21.31%
Commerzbank	998	0.50	23.22%	13.18%	3,162	0.52	20.03%	16.60%
DZ Bank	884	0.67	20.32%	22.24%	2,892	0.58	38.07%	21.91%
Deutsche Bank	1,985	0.57	36.91%	13.45%	2,276	0.35	140.99%	13.79%
Goldman Sachs	5,353	0.26	67.71%	9.31%	5,373	0.26	66.95%	9.27%
Unicredit	-	-	-	-	2,253	0.43	55.90%	12.33%
BNP	2,098	0.68	29.84%	19.25%	3,640	0.72	28.21%	18.78%
Société Générale	1,095	0.54	57.80%	19.28%	315	0.20	12.32%	24.45%
HSBC	-	-	-	-	2,179	0.29	210.05%	13.65%
UBS	1,530	0.49	31.72%	26.03%	722	0.56	24.78%	25.57%
Vontobel	-	-	-	-	3,292	0.52	20.51%	11.43%
Sum	14,297	0.46	45.94%	15.31%	26,373	0.46	61.85%	14.87%

Table 2. *Certificate data separated by the type of bonus certificates. For each issuer, the number of certificates, the average maturity in years, the average moneyness, and the average cushion are given.*

	Minimum	25% Quantile	Median	Mean	75% Quantile	Maximum	Standard deviation
All models	0.00%	0.31%	0.66%	1.11%	1.39%	35.85%	1.32%
Local volatility minus Heston	-34.09%	-0.66%	-0.28%	-0.56%	-0.04%	20.00%	1.07%
Local volatility minus Bates	-35.85%	-1.32%	-0.57%	-0.97%	-0.14%	24.26%	1.36%
Bates minus Heston	-25.07%	-0.07%	0.14%	0.41%	0.66%	20.90%	0.76%
Recalibration	-47.20%	-0.10%	0.00%	0.00%	0.10%	47.19%	0.49%
Heston full minus Heston partial	-20.01%	-0.04%	-0.01%	-0.04%	0.00%	15.10%	0.14%

Table 3. *Key figures of different model risk distributions. The first part refers to model risk based on all models, the second part to pairwise differences between the models, the third part to recalibration risk, and the fourth part to calibration risk for the Heston model (calibrated to all available calibration instruments versus standard options only).*

		Citi	Commerz	DZ	Deutsche	Goldman	UniCredit	BNP	SocGen	HSBC	UBS	Vontobel
Panel A: Classic Bonus Certificates												
Heston	R^2	0.9963*	0.9916*	0.9931*	0.9806*	0.9023	-	0.9886*	0.9636*	-	0.9898*	-
Heston	sd	0.0007*	0.0017*	0.0011*	0.0025*	0.0089	-	0.0014*	0.0058*	-	0.0017*	-
Heston	β	0.9411*	0.9051*	0.9184*	0.7562*	0.6650	-	0.8828*	0.8956*	-	0.8386*	-
Bates	R^2	0.9950	0.9907	0.9857	0.9769	0.9048	-	0.9840	0.9603	-	0.9881	-
Bates	sd	0.0008	0.0018	0.0016	0.0028	0.0088	-	0.0016	0.0061	-	0.0018	-
Bates	β	0.7874	0.8540	0.8597	0.7179	0.6474	-	0.8443	0.8638	-	0.8242	-
Local	R^2	0.9939	0.9855	0.9694	0.9696	0.8531	-	0.9672	0.9286	-	0.9720	-
Local	sd	0.0009	0.0022	0.0024	0.0032	0.0110	-	0.0023	0.0082	-	0.0028	-
Local	β	0.8834	0.8611	0.8372	0.6159	0.5562	-	0.8442	0.8453	-	0.7300	-
Panel B: Capped Bonus Certificates												
Heston	R^2	0.9956*	0.9806*	0.9544*	0.9733*	0.9074	0.9904*	0.9897*	0.9626*	0.9702*	0.6748	0.9699*
Heston	sd	0.0006*	0.0021*	0.0036*	0.0043*	0.0086	0.0033*	0.0012*	0.0017*	0.0058*	0.0095	0.0029*
Heston	β	0.9280*	0.9022*	0.8321*	0.7533*	0.6560	0.9526*	0.8788*	1.0355	0.6780*	0.7992	0.9181*
Bates	R^2	0.9941	0.9797	0.9517	0.9708	0.9101	0.9899	0.9826	0.9443	0.9637	0.6697	0.9683
Bates	sd	0.0007	0.0021	0.0037	0.0045	0.0085	0.0034	0.0016	0.0021	0.0064	0.0096	0.0030
Bates	β	0.8640	0.8659	0.8093	0.7306	0.6392	0.9210	0.8396	0.9677*	0.6543	0.7564	0.8904
Local	R^2	0.9861	0.9614	0.9063	0.9512	0.8629	0.9746	0.9661	0.8986	0.9397	0.6632	0.9407
Local	sd	0.0010	0.0029	0.0052	0.0058	0.0105	0.0054	0.0022	0.0028	0.0082	0.0097	0.0041
Local	β	0.8052	0.8324	0.7512	0.6852	0.5562	0.9060	0.8386	0.8825	0.5733	0.7237	0.8809

Table 4. *Results of the model choice regressions, separated by certificate type (classic vs. capped) and issuer. The adjusted R^2, the standard deviation of the residuals, and the beta coefficient of the theoretical returns are given. The respectively best value is denoted by an asterisk, if the corresponding R^2 exceeds 95%.*