

Empirical Performance of Models for Barrier Option Valuation*

Cathrine Jessen

Dept. of Finance

Copenhagen Business School

e-mail: cj.fi@cbs.dk

Rolf Poulsen[†]

Dept. of Mathematical Sciences

University of Copenhagen

e-mail: rolf@math.ku.dk

August 2011

Abstract

In this paper the empirical performance of five different models for barrier option valuation is investigated: the Black-Scholes model, the constant elasticity of variance model, the Heston stochastic volatility model, the Merton jump-diffusion model, and the infinite activity Variance Gamma model. We use time-series data from the USD/EUR exchange rate market: standard put and call (plain vanilla) option prices and a unique set of observed market values of barrier options. The models are calibrated to plain vanilla option prices, and prediction errors at different horizons for plain vanilla and barrier option values are investigated. For plain vanilla options, the Heston and Merton models have similar and superior performance for prediction horizons up to one week. For barrier options, the continuous-path models (Black-Scholes, constant elasticity of variance, and Heston) do almost equally well, while both models with jumps (Merton and Variance Gamma) perform markedly worse.

JEL classification: G13

Keywords: Barrier option valuation, empirical performance.

*We are deeply indebted to Morten Nalholm for his help cleaning and organizing the barrier option data, and to Fiodar Kilin for his help calibrating the Bates and VG-CIR models.

[†]Corresponding author.

1 Introduction

This paper is an empirical investigation of how well different models work for barrier option valuation. The study is performed using a unique data-set of exchange rate barrier option values.

If we were bold, we would add the qualifiers “first” and “truly” to the word “empirical” in the opening paragraph. That, however, would be pushing the envelope as empirical studies of barrier options are not completely absent from the literature. The performance of “held-until-expiry” hedge portfolios for barrier options on the German DAX index is tested by Maruhn, Nalholm & Fengler (2010), and An & Suo (2009) investigate hedge portfolios for USD/EUR exchange rate barrier options by “marking-to-model”. Actual market values of barrier options are, though, absent from both studies. We know of three previous papers that look at market values of barrier options. Easton, Gerlach, Graham & Tuyl (2004) investigate Australian exchange traded index barrier options, and Wilkens & Stoimenov (2007) study the embedded barrier option in the German Turbo Warrants. But these both work solely in realm of the Black-Scholes model. Carr & Crosby (2010) offer an ingenious model construction that allows for efficient pricing of barrier options but their empirical application is limited to illustrative calibrations for two specific days.

A variety of experimental designs can be used when investigating model performance across time and markets (underlying, *plain vanilla* across strikes and expiry-dates, and exotics). We use one that resembles how the models are used by market participants without violating the basic premise of what constitutes a model: Parametric models are calibrated to liquid plain vanilla options and then used to value exotic options. While this (re-)calibration practice is almost impossible to justify theoretically, a model that does not get the basic contracts about right does not come across as trustworthy when it comes to valuing more advanced products. More specifically, our experimental design is this: On any given day in the sample, say t , each model’s parameters are chosen to obtain the best fit of that day’s plain vanilla option prices across strikes and expiry-dates. We calculate within-that-day (“horizon-0”) pricing errors by comparing observed option values to model values. This is done separately for plain vanilla (“in-sample”) and for barrier options (“out-of-sample” or more tellingly “out-of-market”). We then test the predictive qualities of the models over the horizon h by keeping the time- t calibrated parameters fixed, updating state variables (underlying and possibly volatility) and options (plain vanilla and barrier) to their time- $(t + h)$ values, and registering the discrepancies between model and market values. Some may frown at our use of the word “prediction” and say that we should at least add “conditionally on state variables”, or better yet say that we test “parameter stability”. That is a matter of taste, but what

is not is that such conditional predictions are exactly what matter in a context where plain vanilla options are used as hedge instruments where the idea is to create portfolios that are immunized to changes in state variables.

We work with five popular, yet qualitatively different parametric models: the Black-Scholes model, the constant elasticity of variance model, the Heston stochastic volatility model, the Merton jump-diffusion model, and the infinite activity Variance Gamma model.

For the plain vanilla options we find that the Heston and Merton models have similar performance, and that this performance is superior to the three other models' at horizons of up to five days.

For the barrier options, the performance of the continuous-path models (Black-Scholes, constant elasticity of variance and Heston) is quite similar, and better than those in the few reported previous studies, which all deal with equity markets. And as a general rule, the performance is “half an order of magnitude” worse than for plain vanilla options; more for barrier options whose plain vanilla counterpart is in-the-money when the barrier event happens, less in the opposite case. Both models with jumps (Merton and Variance Gamma) fail miserably for barrier options. These results hold not only at horizon-0 — which could be seen as a self-fulfilling prophesy if market participants use Black-Scholes'ish models for valuation — but for predictions at all horizons.

The rest of the paper is organized as follows: Section 2 describes the data-sets in detail, Section 3 reviews the different models and option pricing techniques, Section 4 reports the results of the empirical analysis, and Section 5 briefly concludes and outlines topics for future research.

2 Data

Our study combines data from two independent sources.¹

Plain vanilla option prices on the USD/EUR exchange rate come from British Bankers' Association.² For each day, we have observations of options with expiries in 1 week, 1 month, 3 months, 6 months, 1 year and 2 years; for the 1 month, 3 months and 1 year expiries we further have prices of options with strikes (roughly) 5% under and 5% over the current exchange rate. Data is given as implied at-the-money forward volatilities,

¹To be entirely precise: three independent sources. To further enhance the data quality, we cross-checked exchange and interest rates against the FED Release H.15.

²This admirable free service was discontinued in early 2008 — possibly not completely surprising as the data quality had deteriorated noticeably throughout 2007. The data we used can be downloaded from <http://www.math.ku.dk/~rolf/papers.html>

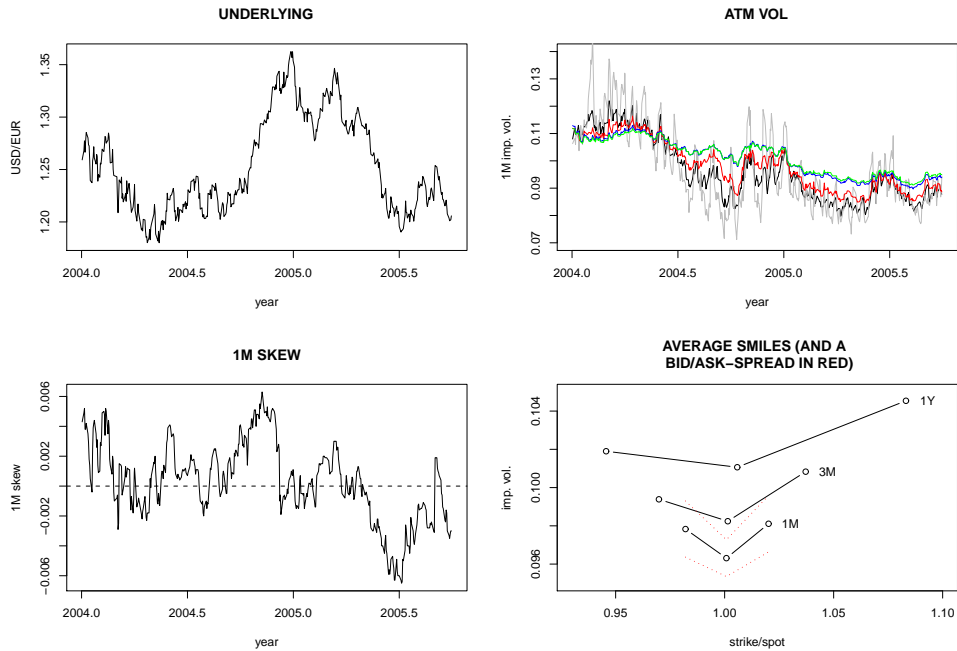


Figure 1: The plain vanilla data from British Bankers’ Association. The top left graph is the USD/EUR exchange rate, i.e. the number of US dollars one has to pay to get 1 Euro. The top right panel is the implied volatility of at-the-money options with one week (gray), one month (black), three months (red), one year (green) and two years (blue) to expiry. The “smiles” in the bottom right panel are the time-averages of implied volatilities (connected) across strikes for expiries of 1 month (lower), 3 months (middle) and 1 year (upper). The dotted red curves around the 1 month smile indicate typical Interbank market bid/ask-spreads. The “skew” depicted in the bottom left graph is the time-series behaviour of the difference between the right and left end-points of the 1 month implied volatility curve.

implied volatilities of 25 delta strangles (average of 25-delta call and 25-delta put minus at-the-money implied volatility) and of 25 delta risk reversals (implied volatility of 25-delta call minus implied volatility of 25-delta put). These implied volatilities can be inverted to give 12 plain vanilla option prices observed each day. Figure 1 shows the data, with option prices being expressed through their implied volatilities. Implied volatility is not constant across time (it decreases throughout our sample), expiry (it increases with time to expiry), or strike (it increases as strike moves away from spot). Compared to equity options, these implied volatilities display a fairly symmetric smile across strikes on average, but there is a randomly varying asymmetry as measured by the skew, i.e. the difference between high- and low-strike implied volatilities.³ The British Bankers’ Association data does not give information about bid/ask-spreads, but accord-

³Others have noticed this and proposed stochastic skew models; Carr & Wu (2007) do it in a Levy setting and in unpublished work Nicole Branger and co-authors use a diffusion framework. We leave the investigation of these models to future work.

	#contracts	#observations	$\bar{\tau}_0$	$\bar{\tau}_i$	$\overline{B/X_0}$	$\overline{K/X_0}$
Total	156	3,108	69	70	–	–
Reverse						
up-and-out call	38	535	44	40	1.041	1.003
up-and-in call	17	309	85	94	1.051	0.997
down-and-out put	16	105	31	25	0.964	1.000
down-and-in put	42	909	69	65	0.954	1.009
Regular						
down-and-out call	33	1,059	93	76	0.969	1.019
down-and-in call	5	52	29	21	0.978	0.991
up-and-out put	5	139	210	177	1.020	0.992
up-and-in put	0	0	–	–	–	–

Table 1: Descriptive statistics of the barrier option data-set; option type, frequency, (calendar) time to expiry, (relative) strike and barrier. A '0' indicates 'at initiation of a particular contract', so the 4th, 6th, and 7th column are averages of contracts at their initiation days of, respectively, time to expiry (τ_0 , measured in calendar days), barrier level (B) relative to spot (X_0), and strike (K) relative to spot. The $\bar{\tau}_i$ denotes time to expiry averaged over all observations (so having $\bar{\tau}_i > \bar{\tau}_0$ is not an error, since there may exist more observations of contracts with relatively high initial time to expiration, τ_0 , thereby contributing to higher average time to expiry, τ_i).

ing to Wystup (2007) a multiplicative spread on volatilities of 1-2% is (or: was at that time) common for at-the-money options in the Interbank market. Or in numerical terms: a typical at-the-money option is sold at 0.101, bought at 0.099.

The exchange rate barrier option data-set⁴ stems from the risk-management department of Danske Bank; the largest Danish bank. Every day the department calls (or: sends a spreadsheet to) the bank's foreign exchange trading desk asking for valuations of all the exchange rate barrier options that the bank currently has on its books. We see no indications in the data that the trading desk is not "doing its job properly"⁵ — such as stale quotes or suspiciously consistent over- or under-valuations of particular trades. The data-set contains values and characteristics for USD/EUR barrier options, but all proprietary information such as counter-party, size and direction of position, and initial price at which the option was sold to (or bought from, but it seems a safe bet that the bank is mostly short in barrier options) the counter-party has been removed.

⁴The data-set plus a detailed description of its construction and organization can be found at <http://www.math.ku.dk/~rolf/papers.html>

⁵The life-span of the barrier options is much shorter than the "bonus horizon"; that should alleviate moral hazard.

This would of course make for interesting reading and research, but on the plus-side the lack of sensitive information means that the data has been released for research without “strings attached”. More specifically, the data-set consists of daily observations of continuously monitored zero-rebate barrier contracts covering the period January 2, 2004 to September 27, 2005. We consider only single barrier options and disregard options with values lower than 10^{-5} and/or less than 7 days to expiry (thus staying within the expiry-range of the plain vanilla calibration instruments). This leaves us with a total of 3,108 observations on 156 individual contracts. These are broken down by characteristics in Table 1. We see that the strikes are mostly set very close to the exchange-rate at contract initiation (i.e. the corresponding plain vanilla option is at-the-money), that the typical time to expiry is 70 (calendar) days, and that it is common to have the barrier 2-5% away from initial spot. This means that the barrier options fit nicely into the range (expiry and moneyness-wise) of our plain vanilla calibration instruments. Another important feature of a barrier option is whether it is *reverse* (also known as *live-out*) or not, which we term *regular*. A barrier option is of reverse type if the corresponding plain vanilla option is in-the-money when the barrier event (knock-in or knock-out) happens. This means that values of reverse barrier options change very rapidly in the vicinity of the barrier: There is a big difference between just crossing, and not crossing; exploding Greeks and gap risk are other terms used to describe this phenomenon. This makes them hard to hedge — be that statically or dynamically, see Nalholm & Poulsen (2006*a*, Table 2) for instance. The reverse barrier options are the down-and-out put and the up-and-out call and their knock-in counterparts. From Table 1 we see that the data-set is fairly balanced; in general reverse-type options are more common (73% of contracts, 59% of observations), but the single-most observed contract is the (regular) down-and-out call.

A final sanity check of the barrier option data is given in Figure 2. It shows (all) the barrier options’ “implied” volatilities as expressed by the bank.⁶ We see that barrier option volatilities line up reasonably closely to the implied volatilities of the plain vanilla options.

⁶The reason for the quotes and the disclaimer is that even in the Black-Scholes model barrier option values are not monotone functions of volatility. Hence, given an observed barrier option value there may be multiple sensible input volatilities that match the observation, and thus implied volatility is not uniquely defined. The bank’s data-set contains both “implied” volatilities and actual prices; we use the former only for graphical purposes.




Figure 2: The barrier option data from Danske Bank. Each gray circle represents a barrier option data point in terms of an “implied” on the spot volatility. The fully drawn curve is the implied volatility of the 1 month, at-the-money plain vanilla option.

3 Model selection and pricing methods

We consider five alternative models for the exchange rate. The model selection aims at including models with different features: a model with state dependent volatility versus one with stochastic volatility and a model with low jump activity and large jumps versus one with high jump activity and small jumps. One requirement though is the existence of reasonable methods for pricing barrier options either analytically or numerically. Under these criteria we have chosen the following models, which are frequently encountered in the literature: the constant elasticity of variance model (CEV), the stochastic volatility model of Heston, the Merton jump-diffusion model and the infinite activity Variance Gamma model (VG). Our benchmark model is that of Black-Scholes. All models have the Black-Scholes model as a special or limiting case, but apart from that they are as non-nested as can be, thus covering a large range of qualitatively different (and popular) models. The models and pricing methods “at a glance” are shown in Table 2; more detailed descriptions are given in the following subsections.

Model	Plain Vanilla	Barriers
Black-Scholes	Closed-form.	All-encompassing closed-form formulation in Rubinstein & Reiner (1991).
Constant elasticity of variance	Closed-form à la Schroder (1989) w/ Ding's algorithm for the non-central χ^2 -distribution.	Collocation à la Nalholm & Poulsen (2006 <i>b</i>). Inversion à la Davydov & Linetsky (2001).
Heston stochastic volatility	Fourier inversion using the formulation in Lipton (2002).	Simulation à la Andersen (2008) w/ bridge and control. PDE solution à la Foulson & In 't Hout (2010). Perturbation expansion à la Wong & Chan (2008).
Merton jump-diffusion	The original formula by Merton (1976) is found more efficient than Fourier inversion.	Simulation à la Joshi & Leung (2007) and Metwally & Atiya (2002).
Variance Gamma	Fourier inversion w/ tricks from Lee (2004) and Glasserman (2004).	Simulation à la Glasserman (2004) w/ tricks from Avramidis (2004).

Table 2: Annotated taxonomy of pricing.

3.1 The Black-Scholes model

In the Black-Scholes model the foreign exchange rate X follows a geometric Brownian motion under the risk-neutral pricing measure:⁷

$$dX_t = (r_d - r_f)X_t dt + \sigma X_t dW_t,$$

where r_d and r_f denote the assumed-constant domestic (US) and foreign (Euro) interest rates. In this setup, closed-form formulas for both plain vanilla and barrier option prices exist and will be used for pricing. It is well-known that the one-parameter Black-Scholes model is not the best model to describe observed option prices — especially not for a wider range of strikes and maturities simultaneously. However, it may still turn out to be the preferred model choice for pricing barrier options due to its fast, stable and easily implementable pricing procedure.

⁷By construction the price calibration estimates the pricing measure used by the market, thus all parameters are under the/a risk-neutral pricing measure; say Q . For our price analysis, this is not a restriction; if we were to study construction and performance of hedge portfolios both the risk-neutral pricing measure and the real-world measure would matter — though possibly less so in practice than in theory, see Poulsen, Schenk-Hoppé & Ewald (2009) and Siven & Poulsen (2009, Table 4).

3.2 The Constant elasticity of variance model

A minimal extension of the Black-Scholes model is the constant elasticity of variance model, in which the foreign exchange rate has the risk neutral dynamics

$$dX_t = (r_d - r_f)X_t dt + \sigma X_t^\alpha dW_t,$$

where α denotes the so-called elasticity of variance. In this model there are two parameters, α and σ , to be estimated. For $\alpha < 1$ volatility increases as the exchange rate falls; vice versa for $\alpha > 1$.

For pricing plain vanilla options we use the closed-form formula of Schroder (1989). Several methods for pricing barrier options exist. Numerical techniques such as the finite difference method and Monte Carlo simulation can be applied. Alternatively, as we have chosen to do here, barrier option prices may be found via collocation as demonstrated in Nalholm & Poulsen (2006b). Analytical formulas for barrier option prices based on inversion techniques do exist (see Davydov & Linetsky (2001)), however, these are rather involved and in our experience there is no real gain with respect to computation time compared to the direct numerical approaches.

3.3 Heston's stochastic volatility model

For the stochastic volatility model we have chosen the Heston model, where the exchange rate and its instantaneous variance follow

$$\begin{aligned} dX_t &= (r_d - r_f)X_t dt + \sqrt{v_t}X_t dW_t^1, \\ dv_t &= \kappa(\theta - v_t)dt + \eta\sqrt{v_t}dW_t^2. \end{aligned}$$

Here θ is the long term level of variance, κ is the speed of mean reversion, η is referred to as the volatility of volatility, and the driving Brownian motions have correlation ρ , leading to a skew in implied volatilities. In the Heston model there are four parameters, κ , θ , η and ρ , plus one state variable, v , to be estimated. The (conditional) characteristic function of X can be found in closed form; this was first done in Heston (1993). This means that plain vanilla option pricing becomes a question of one-dimensional numerical integration; inverting a transform. There is a sizable literature on this, see Lee (2004). We prefer the quadratic denominator formula of Lipton (2002). Closed-form solutions for barrier options exist (see Lipton (2001)) in the case where domestic and foreign short rates are equal and correlation is zero, but Faulhaber (2002) shows that there is no simple way to relax those assumptions, which are unrealistic to impose on our data. During the sample period the US short rate r_d decreases from approximately 4% to 1%, while the European short rate r_f is more or less constant at 2.1%, and since we do see

an implied skew in our data, fixing $\rho = 0$ is also too restrictive. Alternatives are the PDE method described by Foulson & In 't Hout (2010) or Monte Carlo simulation using a quadratic exponential discretization scheme for the volatility process developed by Andersen (2008). We use the latter and combine it with a Black-Scholes model control variate. The simulation is run with 10 000 trajectories, each with time steps of one day.

3.4 Merton's jump-diffusion model

The Black-Scholes model can also be extended to include jumps in the exchange rate as done by Merton:

$$dX_t = (r_d - r_f - \lambda \mathbf{E}^Q(Z_t - 1))X_{t-}dt + \sigma X_{t-}dW_t + X_{t-}(Z_t - 1)dN_t,$$

where N is a Poisson process with intensity λ , and $\log Z_t \sim \mathcal{N}(\mu_Z, \sigma_Z)$ describes the relative jump size as being normally distributed with mean μ_Z and variance σ_Z . The Merton model has four parameters to be estimated: σ , μ_Z , σ_Z and λ .

Pricing plain vanilla options in this model can be done by Fourier inversion techniques or — in our experience more efficiently — by using the original formula from Merton (1976). Barrier option prices are found by Monte Carlo simulation methods as suggested by Metwally & Atiya (2002) and Joshi & Leung (2007) via the use of importance sampling. Again we use 10 000 trajectories with daily sampling in the simulation, and check that this is indeed sufficient to generate stable barrier prices.

3.5 The Variance Gamma model

Another class of jump models are models exhibiting infinite jump activity as e.g. the Variance Gamma (VG) model proposed by Madan & Seneta (1990). The foreign exchange rate under the risk neutral measure in the VG model is of the form

$$X_t = X_0 \exp\{(r_d - r_f)t + Y_t^{\text{VG}} + \omega t\},$$

where $Y_t^{\text{VG}} = \theta G_t^\nu + \sigma W(G_t^\nu)$ is a variance gamma process; a time changed Brownian motion with drift, $\theta t + \sigma W_t$, using a gamma process G^ν with volatility ν as the stochastic clock. The martingale correction term $\omega = \frac{1}{\nu} \ln(1 - \theta\nu - \frac{1}{2}\sigma^2\nu)$ ensures that the expected rate of return on assets equals the risk-neutral rate $r_d - r_f$. The parameter ν controls for excess kurtosis and θ for skewness. The limit when $\nu \rightarrow 0$ (in which case the influence from θ also disappears) is the Black-Scholes model. Figure 3 shows some simulated paths of VG processes, and illustrates that for small ν -values, the process looks diffusion'ish, while high ν gives a more Poisson-jump-like appearance.

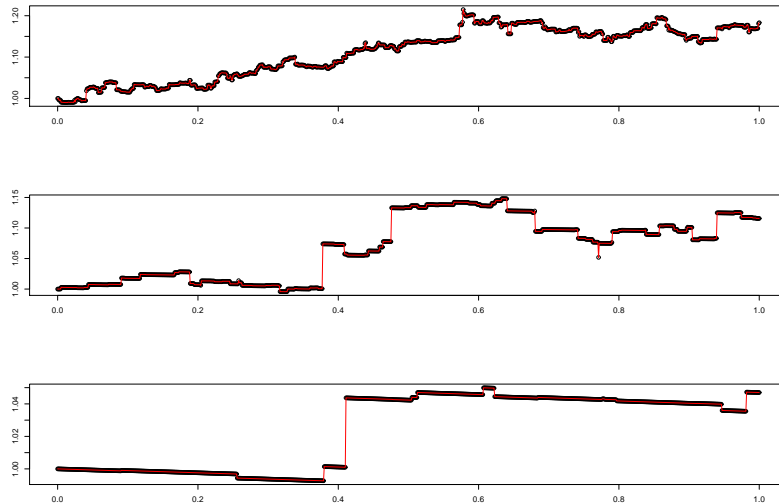


Figure 3: Simulated paths of Variance Gamma processes. The parameter values are $\sigma = 0.1$, $\theta = 0.0085$, and $\nu = 0.01$ (top), 0.1 (middle), 1 (bottom).

Plain vanilla option prices can be found by Fourier inversion, as done in e.g. Lee (2004). Barrier option prices can be found by simulation methods as presented in Glasserman (2004) or the double-gamma bridge sampling algorithm by Avramidis (2004). 10 000 trajectories with daily sampling are used to generate the barrier prices.

4 Empirical results

4.1 Calibration and plain vanilla option valuation

On any specific date (t) and for any model j (naturally indexed by $\{\text{BS}, \text{CEV}, \text{H}, \text{M}, \text{VG}\}$), we estimate the parameter, $\vartheta_j(t)$, by minimizing the sum of absolute differences between the observed implied volatilities (IV) and the model's implied volatilities. Or with symbols:

$$\widehat{\vartheta}_j(t) = \arg \min_{\vartheta} \sum_{i|t(i)=t} |IV^{\text{obs}}(i) - IV^{\text{model } j}(i; \vartheta)|,$$

where the notational philosophy is that i denotes observations, and $t(\cdot)$ maps an observation to its date.

Implied volatilities put option prices in a comparable scale across strikes and expiries. Minimizing differences to raw prices does not alter our results but makes the numbers harder to relate to. One could also minimize differences to relative prices but in our

experience that tends to put too much weight on out-of-the-money options.⁸

Sample characteristics of estimators are given in Table 3. The calibrated parameters are not constant over the sample period, but they are more stable than the meta-analysis for S&P500 that is reported in Gatheral (2006, Table 5.4) indicating that exchange rate markets are more benign than equity markets. Only the dangerously naive observer would claim that a price outside the bid/ask-spread is an arbitrage opportunity, but it is nonetheless a sensible yardstick. And on that count the models separate into two categories; Merton and Heston hit about two-thirds of bid/ask-spreads, and the rest about 40%. Other remarks:

- The difference between the average instantaneous variance, v_0 , and the Heston model's (risk-adjusted parameter) θ reflects the typically increasing term structure of volatility.
- In the Merton model most (75%) of the variance of daily returns is caused by the diffusion component.
- The sample mean of the VG estimates were used to generate the middle path in Figure 3; the paths have a visible, but not extreme, non-diffusive character.

⁸To illustrate: An implied volatility difference of 0.001 gives a 1% relative price difference for the at-the-money option, but 2% for the out-of-the-money option. A variety of weighting schemes have been suggested in the literature; *to every man his own*. An interesting point is made by Cont & Tankov (2004, p. 439) that calibration to squared implied volatility differences corresponds (to a first-order approximation) to using a vega-weighted average of price differences.

Model	Black-Scholes		CEV		Heston				
	σ		σ	α	v_0	θ	κ	η	ρ
mean	0.0998		0.0891	1.50	0.00964	0.0120	2.22	0.183	0.0634
std.dev	0.0088		0.0147	0.597	0.00244	0.00139	0.525	0.0633	0.117
median	0.0996		0.0899	1.24	0.00889	0.0125	2.12	0.184	0.0731
min	0.0877		0.0451	-0.091	0.00615	0.00894	0.879	0.00186	-0.465
max	0.116		0.118	3.39	0.0161	0.0150	7.91	0.412	0.470
within bid/ask	40.3 %		38.8%		69.2%				

Model	Merton				VG		
	σ	λ	σ_Z	μ_Z	σ	ν	θ
mean	0.0851	1.24	0.0443	0.00339	0.105	0.102	0.00847
std.dev	0.0157	0.104	0.0208	0.0116	0.00387	0.0851	0.0666
median	0.0806	1.26	0.0510	0.00280	0.105	0.0898	0.00929
min	0.0593	0.587	$5.00 \cdot 10^{-5}$	-0.0325	0.0921	$1.08 \cdot 10^{-4}$	-0.310
max	0.118	1.54	0.077	0.0524	0.113	0.357	0.294
within bid/ask	65.6%				44.3%		

Table 3: Sample characteristics of parameter estimates. The “within bid/ask” rows give the percentage of the calibrated model prices that fall within a bid/ask spread of ± 0.0015 on implied volatility.




Figure 4: Prediction errors for plain vanilla options at different horizons.

A particular model may perform well on the data that it is calibrated to, but have poor predictive qualities (think of fitting a high-order polynomial to “a regular line with noise”). To investigate this, Figure 4 shows the five models’ prediction errors at different (business day) horizons. The fully drawn, differently colored curves show the models’ average absolute implied volatility differences for increasing horizons, i.e.

$$\frac{1}{\#\text{obs. dates} - h} \sum_t \frac{1}{\#i|t(i) = t} \sum_{i|t(i)=t} |IV^{\text{obs}}(i) - IV^{\text{model } j}(i; \hat{\vartheta}_{t(i)-h})|.$$

It is only for the Variance Gamma model that the ordering is changed when we look at predictions; it is (slightly) better than Black-Scholes at horizon 0, but worse at longer horizons. The errors of the Merton model are marginally lower than those of the Heston model at all horizons, and the differences are statistically significant (at a 5% level) at horizons of two days or more. As a rule-of-thumb-quantification of how much better the Heston and Merton models perform, we can look at the horizons where their prediction errors match the horizon-0 errors of the others models, i.e. the points, say h_j , on the abscissa where the red and blue curves cross the dash-dotted horizontal lines. A way to

interpret these numbers is to say: “Using model j with ‘freshly estimated’ parameters is (on average) as good as using a Heston or Merton model with h_j day old parameters”. We see that the Heston and Merton models are caught up with by the other models after about one week.

4.1.1 Combining stochastic volatility and jumps; the Bates model

One may suspect a combination of the Heston and Merton models to perform even better. A stochastic volatility model including Poisson jumps in the exchange rate — also known as a Bates model following Bates (1996) — has the dynamics

$$\begin{aligned} dX_t &= (r_d - r_f - \lambda \mathbf{E}(Z_t - 1))X_{t-}dt + \sqrt{v_t}X_{t-}dW_t^1 + X_{t-}(Z_t - 1)dN_t \\ dv_t &= \kappa(\theta - v_t)dt + \eta\sqrt{v_t}dW_t^2, \end{aligned}$$

with $\rho dt = \text{cor}(dW_t^1, dW_t^2)$. The Bates model has a total of seven parameters, κ , θ , η , ρ , λ , μ_Z and σ_Z , plus one state variable, v_0 . Calibration of this model is a numerically delicate matter but can be carried out as suggested by Kilin (2011).

For the plain vanilla data the Bates model’s average absolute implied volatility error is up to two significant digits (and no statistical significance) identical to the Merton and Heston models’. Since the Bates model does not improve the results over Merton and Heston models, we see no reason to implement further extensions of the model, like e.g. the Universal Volatility model introduced in Lipton (2002) and used for barrier option pricing in Lipton & McGhee (2002), in a setting with correlation $\rho = 0$.

4.1.2 Combining Levy-models and stochastic volatility; the VG-CIR model

A way to introduce stochastic volatility into pure jump models such as the Variance Gamma is to subject the driving process to a random time-change, i.e. to work with $Y_{Z_t}^{\text{VG}}$ where Z_t is an increasing stochastic process. Carr, Geman, Madan & Yor (2003) show how characteristic functions in some cases can be expressed by composition of the Laplace-transform of the time change process and the characteristic function of the original model. A convenient choice of time-change process is an integrated Cox-Ingersoll-Ross process (independent of the original Y^{VG} -process), whose Laplace-transform is part of the interest rate theory vocabulary. Again, calibration of this six-parameter (and one more or less latent state variable; the current value of the subordinator) VG-CIR model is a delicate matter for which we refer to Kilin (2011).

The cross-sectional average implied volatility error for the model is 0.23%; lower than the VG model’s error, but not as good as the Heston and Merton models.⁹

⁹Detailed results for the Bates and VG-CIR, corresponding to those in table 4 and 5, will be provided upon request.

% h split		Black-Scholes		CEV		Heston	
		$\overline{e^{BS}}$	$ \overline{e^{BS}} $	$\overline{e^{CEV}}$	$ \overline{e^{CEV}} $	$\overline{e^H}$	$ \overline{e^H} $
0	all	0.102 [‡]	4.93	-0.0028 [‡]	4.49	0.53	2.21
	(ATM , OTM)	(1.15 , -0.495)	(3.48 , 6.39)	(1.03 , -1.03)	(3.38 , 5.60)	(1.95 , -0.89)	(2.62 , 1.80)
1	all	0.184 [‡]	5.19	0.081 [‡]	4.86	0.618	3.02
	(ATM , OTM)	(1.20 , -0.831)	(3.67 , 6.71)	(1.08 , -0.916)	(3.59 , 6.14)	(2.01 , -0.769)	(3.03 , 3.02)
5	all	0.537	6.06	0.479	5.84	0.977	4.47
	(ATM , OTM)	(1.42 , -0.345 [‡])	(4.26 , 7.85)	(1.33 , -0.369 [‡])	(4.18 , 7.50)	(2.23 , -0.275 [‡])	(3.87 , 5.07)

% h split		Merton		VG	
		$\overline{e^M}$	$ \overline{e^M} $	$\overline{e^{VG}}$	$ \overline{e^{VG}} $
0	all	-0.348	2.31**	0.58	4.79
	(ATM , OTM)	(-0.36 , -0.34)	(1.80 , 2.83)	(1.07 , 0.085 [‡])	(2.95 , 6.64)
1	all	-0.257	3.00**	0.590	5.40
	(ATM , OTM)	(-0.30 , -0.21 [†])	(2.31 , 3.69)	(1.08 , 0.0998 [‡])	(3.41 , 7.39)
5	all	0.118 [‡]	4.30	0.675	6.51
	(ATM , OTM)	(-0.064 [‡] , 0.30 [†])	(3.35 , 5.25*)	(1.15 , 0.201 [‡])	(4.35 , 8.68)

Table 4: The table reports average relative pricing errors and average absolute relative pricing errors (in %) for plain vanilla options in the the various calibrated models. The relative pricing errors are pooled and divided by moneyness; and at different prediction horizons, h . For the raw errors, daggers show statistically significant differences from 0 at 1% or 5% levels. For the numerical errors we use a “reverse” significance convention: Star(s) next to a number indicates that a paired “equal to Heston?”-test cannot be rejected at the 5% (single star) or 1% (double star) level.

%	h	split	Black-Scholes		CEV		Heston	
			$\overline{\xi^{BS}}$	$ \overline{\xi^{BS}} $	$\overline{\xi^{CEV}}$	$ \overline{\xi^{CEV}} $	$\overline{\xi^H}$	$ \overline{\xi^H} $
0	all		2.71	8.27**	0.13 [‡]	8.50	-3.48	7.82
	(rev. , str.)		(4.12 , 0.61 [†])	(11.1** , 4.0**)	(-0.37 [‡] , 0.88)	(11.1** , 4.68)	(-4.61 , -1.80)	(10.4 , 3.89)
	(up , down)		(-0.37 [‡] , 4.13)	(6.35 , 9.16)	(-1.47 , 0.87 [†])	(7.66** , 8.94)	(-4.09 , -3.20)	(7.83 , 7.81)
1	all		2.75	8.33**	0.075 [‡]	8.69**	-3.09	8.22
	(rev. , str.)		(4.18 , 0.61 [†])	(11.2** , 4.02**)	(-0.55 [‡] , 1.00)	(11.4* , 4.59)	(-4.61 , -1.68)	(10.5 , 3.78)
	(up , down)		(-0.50 [‡] , 4.25)	(6.33 , 9.26*)	(-1.75 , 0.92 [‡])	(7.63** , 9.19**)	(-4.21 , -2.57)	(7.84 , 8.39)
5	all		3.60	10.6**	0.91	10.3**	-2.31	10.1
	(rev. , str.)		(5.16 , 1.23)	(14.0** , 5.49)	(0.30 [‡] , 1.83)	(13.2** , 6.02)	(-3.03 , -1.23)	(13.6 , 4.83)
	(up , down)		(-0.010 [‡] , 5.26)	(8.38 , 11.6)	(-1.39 , 1.98)	(9.36** , 10.8**)	(-3.87 , -1.59)	(9.50 , 10.3)

17

%	h	split	Merton		VG	
			$\overline{\xi^M}$	$ \overline{\xi^M} $	$\overline{\xi^{VG}}$	$ \overline{\xi^{VG}} $
0	all		-22.5	29.3	66.7	79.5
	(rev. , str.)		(-21.1 , -24.7)	(31.7 , 25.7)	(106 , 8.39)	(123 , 14.4)
	(up , down)		(-8.91 , -28.8)	(26.3 , 30.7)	(49.6 , 74.7)	(65.5 , 86.0)
1	all		-22.5	29.3	63.4	76.0
	(rev. , str.)		(-21.1 , -24.7)	(31.7 , 25.6)	(100 , 8.39)	(118 , 14.1)
	(up , down)		(-9.01 , -28.7)	(27.0 , 30.8)	(48.1 , 70.5)	(63.6 , 81.7)
5	all		-21.3	30.3	68.0	78.1
	(rev. , str.)		(-19.6 , -23.9)	(33.5 , 25.6)	(102 , 10.1)	(119 , 15.9)
	(up , down)		(-7.82 , -27.6)	(29.3 , 30.8)	(44.5 , 75.1)	(60.4 , 86.3)

Table 5: Relative pricing errors for barrier options in %. We use notation and significance signs as in Table 4.

4.2 Barrier option valuation

We now turn to the main question: How well do the different models perform when it comes to valuing barrier options? Implied volatilities for barrier options are not well-defined and their raw prices differ several orders of magnitude. Therefore we will report relative price errors for the barrier options. To fix notation, the (percentage) relative error for the i 'th observation for the j 'th model at horizon h is

$$\xi_{i,h}^j = 100 \times \frac{B^{\text{model } j}(i; \hat{\vartheta}_j(t(i) - h)) - B^{\text{obs}}(i)}{B^{\text{obs}}(i)},$$

where the notational philosophy is as before, and the $B(i; \cdot)$'s denote values of barrier options with the appropriate characteristics. For comparison we report relative errors of plain vanilla options too. These are defined analogously and denoted by $e_{i,h}^j$.

For plain vanilla options, the estimation procedure ensures that each model's average errors are small.¹⁰ Thus plain vanilla comparisons should be made based on some measure of dispersion such as standard or mean absolute deviation. For barrier options both average errors and their dispersion are relevant measures of a model's quality. Therefore our tables (Table 5 for barrier options, Table 4 for plain vanilla) report sample averages of both errors and absolute errors. For the errors, † and ‡ indicate that there is no significant difference from 0 at, respectively, the 5% and 1% levels. For the absolute errors, * and ** indicate that errors are not significantly different from those of the Heston model; this is the result of a paired test based on absolute error differences.

In Table 4 we have sub-divided the relative plain vanilla option pricing errors into out-of-the-money (OTM) and at-the-money (ATM) errors. General for all five models (except Heston's errors at horizon 0), is that they produce larger relative pricing errors for options out-of-the-money. This is in line with the previous observation that a given implied volatility error corresponds to a larger relative price error for an out-of-the-money option than for an option at-the-money.

From Table 5 we see that the three continuous-path models (Black-Scholes, CEV and Heston) have quite similar behaviour when it comes to barrier option valuation. The CEV model is most accurate with regards to average price errors (0.1% vs. -3.5% for Heston and 2.7% for Black-Scholes), while the Heston model has the lowest dispersion (average absolute error of 7.8% vs. 8.2% for Black-Scholes and 8.5% for CEV). One could explain the good behaviour of the Black-Scholes model as a self-fulfilling prophecy; market participants use Black-Scholes formulas because that is what is on their computers. But if that were the only reason, we would expect to see rapid deterioration

¹⁰We minimize averages of absolute differences of implied volatilities. Therefore average errors (raw and particularly relative) are not *exactly* zero.

in the Black-Scholes model’s predictive quality.¹¹ We do not; conclusions are invariant to the prediction horizon.

Looking at the plain vanilla benchmark in Table 4 we see that the continuous-path models’ error dispersions for barrier options are two to five times larger than for plain vanilla options. Or differently put, barrier options are half an order of magnitude harder to price.¹² For the jump-models (Merton and Variance Gamma), the story is quite different. The Merton model is bad (dispersion of 29%; about four times that of the continuous-path models), and the Variance Gamma model is worse (dispersion of 79%; a ten-fold increase). Again, this holds at all horizons. One could argue that “that is because continuous-path models systematically underestimate knock-out probabilities. You should still use models with jumps.” But that fails to explain why the Merton model undervalues the barrier options (by 23% on average) and the Variance Gamma model overvalues them (by 67% on average). For plain vanilla options, the Merton model was arguably one of the best performing models, and Variance Gamma was on par with Black-Scholes and CEV. Thus their poor ability to explain barrier option values again emphasizes the model risk aspect pointed out by Schoutens, Simons & Tistear (2004), Hirta, Courtadon & Madan (2002), Detlefsen & Härdle (2007) and numerous other papers: Models may produce very similar prices of plain vanilla options yet differ markedly for exotic options.

To understand the models’ pricing performance for barrier options we have analyzed inter-model differences; we simply changed “observed values” to “Black-Scholes values” in the definition of relative errors. This reveals that the Black-Scholes and CEV model values typically are closer to each other, than they each are to data (the average absolute CEV-to-Black-Scholes error is 4.8% compared to about 8% for each model with observed values as reference point), while the average absolute Heston-to-Black-Scholes error is 8.4%.

To further detect patterns, we have sub-divided errors according to different criteria: reverse vs. regular and up vs. down. Results are also reported in Table 5; the numbers in parentheses. First, we see that error dispersions are markedly larger for reverse barrier options than for the regular ones; sample averages of absolute errors are 2-3 times higher.

¹¹This line of thought is equally valid with Black-Scholes substituted by other named models.

¹²Since the plain vanilla option data and the barrier option data stem from two different sources, we have no guarantee that the data sets are collected at the same time of day and this may influence the models’ performance with respect to barrier option pricing. To check whether such time asynchronicity affects our results, we have recalculated the barrier option errors using parameters estimated using plain vanilla prices one day prior and one day after the barrier option observation. Both experiments results in marginally worse barrier option errors. Therefore, if a time asynchronicity exists, it has no significant influence on our results.

Given the difficulties in hedging the reverse options due to their exploding Greeks, this increased dispersion between market and model values may be understandable, but it should be noted that there is no clear pattern for the average errors from the reverse-regular stratification. There is little effect from the up-down split which we interpret as more evidence that exchange rate markets are reasonably symmetric.

With respect to barrier option pricing, the Bates model performs better than the Merton and VG models but significantly worse than the Heston model with average absolute errors of 24.7% ($h = 0$) and 25.6% ($h = 5$). So adding jumps to the Heston model merely worsens the model's barrier option valuation abilities. The same is true when introducing stochastic volatility into the VG model as in the VG-CIR model.

5 Conclusion

We investigated empirical barrier option values, and found that in general the continuous-path models did equally well in explaining the market data, while models with jumps turned out to be quite inaccurate, this despite the jump-diffusion model being — arguably — the best performing model for plain vanilla options.

A logical next step is to investigate how well the barrier options can be hedged, dynamically, statically or by some hybrid hereof. A particularly interesting question, that the barrier option data-set allows us to shed (some) light on, is the benefit of applying a portfolio — rather than “each option on its own” — approach to hedging.

References

- An, Y. & Suo, W. (2009), ‘An empirical comparison of option pricing models in hedging exotic options’, *Financial Management* **38**, 889–914.
- Andersen, L. (2008), ‘Simple and efficient simulation of the Heston stochastic volatility model’, *Journal of Computational Finance* **11**, 1–42.
- Avramidis, A. N. (2004), Efficient Pricing of Barrier Options with the Variance-Gamma Model, in ‘Proceedings of the 2004 Winter Simulation Conference’, IEEE press, pp. 574–1578.
- Bates, D. (1996), ‘Jumps and Stochastic Volatility: The Exchange Rate Processes implicit in Deutsche Mark Options’, *Review of Financial Studies* **9**, 69–107.
- Carr, P. & Crosby, J. (2010), ‘A class of Levy Process Models with almost exact calibration to both barrier and vanilla FX options’, *Quantitative Finance* **10**, 1115–1136.

- Carr, P., Geman, H., Madan, D. B. & Yor, M. (2003), ‘Stochastic Volatility for Levy Processes’, *Mathematical Finance* **13**, 345–382.
- Carr, P. & Wu, L. (2007), ‘Stochastic Skew for Currency Options’, *Journal of Financial Economics* **86**, 213–247.
- Cont, R. & Tankov, P. (2004), *Financial Modelling With Jump Processes*, Chapman & Hall.
- Davydov, D. & Linetsky, V. (2001), ‘Pricing and Hedging Path-Dependent Options Under the CEV Process’, *Management Science* **47**, 949–965.
- Detlefsen, K. & Härdle, W. (2007), ‘Calibration Risk for Exotic Options’, *Journal of Derivatives* **14**, 47–63.
- Easton, S., Gerlach, R., Graham, M. & Tuyl, F. (2004), ‘An Empirical Examination of the Pricing of Exchange-Traded Barrier Options’, *Journal of Futures Markets* **24**, 1049–1064.
- Faulhaber, O. (2002), Analytic methods for pricing double barrier options in the presence of stochastic volatility, Master’s thesis, Diploma Thesis, University of Kaiserslautern, <http://www.oliver-faulhaber.de/mathematik/thesis>.
- Foulson, S. & In ’t Hout, K. J. (2010), ‘ADI Finite Difference Schemes for Option Pricing in the Heston Model with Correlation’, *International Journal of Numerical Analysis and Modeling* **7**, 303–320.
- Gatheral, J. (2006), *The Volatility Surface*, Wiley Finance.
- Glasserman, P. (2004), *Monte Carlo Methods in Financial Engineering*, Springer.
- Heston, S. (1993), ‘A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options’, *Review of Financial Studies* **6**, 327–344.
- Hirsa, A., Courtadon, G. & Madan, D. B. (2002), ‘The Effect of Model Risk on the Valuation of Barrier Options’, *Journal of Risk Finance* **4**, 47–55.
- Joshi, M. & Leung, T. (2007), ‘Using Monte Carlo Simulation and Importance Sampling to Rapidly Obtain Jump-Diffusion Prices of Continuous Barrier Options’, *Journal of Computational Finance* **10**, 93–105.
- Kilin, F. (2011), ‘Accelerating the Calibration of Stochastic Volatility Models’, *The Journal of Derivatives* **18**, 7–16.

- Lee, R. (2004), ‘Option Pricing by Transform Methods: Extensions, Unification, and Error Control’, *Journal of Computational Finance* **7**, 51–86.
- Lipton, A. (2001), *Mathematical Methods For Foreign Exchange: A Financial Engineer’s Approach*, World Scientific.
- Lipton, A. (2002), ‘The vol smile problem’, *Risk Magazine* **15**(February), 61–65.
- Lipton, A. & McGhee, W. (2002), ‘Universal barriers’, *Risk Magazine* **15**(May), 81–85.
- Madan, D. & Seneta, E. (1990), ‘The Variance Gamma (V.G.) Model for Share Market Returns’, *Journal of Business* **63**, 511–524.
- Maruhn, J., Nalholm, M. & Fengler, M. (2010), ‘Static Hedges for Reverse Barrier Options with Robustness Against Skew Risk: An Empirical Analysis’, *Quantitative Finance* **11**, 711–727.
- Merton, R. (1976), ‘Option pricing when the underlying stock returns are discontinuous’, *Journal of Financial Economics* **5**, 125–144.
- Metwally, S. A. K. & Atiya, A. F. (2002), ‘Using Brownian Bridge for Fast Simulation of Jump-Diffusion Processes and Barrier Options’, *Journal of Derivatives* **10**, 43–54.
- Nalholm, M. & Poulsen, R. (2006a), ‘Static Hedging and Model Risk for Barrier Options’, *Journal of Futures Markets* **26**, 449–463.
- Nalholm, M. & Poulsen, R. (2006b), ‘Static hedging of barrier options under general asset dynamics: Unification and application’, *Journal of Derivatives* **13**, 46–60.
- Poulsen, R., Schenk-Hoppé, K. R. & Ewald, C.-O. (2009), ‘Risk Minimization in Stochastic Volatility Models: Model Risk and Empirical Performance’, *Quantitative Finance* **9**, 693–704.
- Rubinstein, M. & Reiner, E. (1991), ‘Breaking Down the Barriers’, *Risk Magazine* **4**(September), 28–35.
- Schoutens, W., Simons, E. & Tistart, J. (2004), ‘A Perfect Calibration! Now What?’, *Wilmott Magazine* (10), 66–78.
- Schroder, M. (1989), ‘Computing the Constant Elasticity of Variance Option Pricing Formula’, *Journal of Finance* **44**, 211–219.
- Siven, J. & Poulsen, R. (2009), ‘Auto-Static for the People: Risk-Minimizing Hedges of Barrier Options’, *Review of Derivatives Research* **12**, 193–211.

- Wilkins, S. & Stoimenov, P. A. (2007), 'The pricing of leverage products: An empirical investigation of the German market for 'long' and 'short' stock index certificates', *Journal of Banking and Finance* **31**, 735–750.
- Wong, H. Y. & Chan, C. M. (2008), 'Turbo warrants under stochastic volatility', *Quantitative Finance* **8**, 749–751.
- Wystup, U. (2007), *FX Options and Structured Products*, John Wiley & Sons.