

Equity with Markov-modulated dividends

GIUSEPPE DI GRAZIANO*
Statistical Laboratory
University of Cambridge

L.C.G ROGERS†
Statistical Laboratory
University of Cambridge

16th December 2005

Abstract

We introduce a simple model for the pricing of European style options when the underlying dividend process is given by a geometric Brownian motion with Markov-modulated coefficients. It turns out that the corresponding stock process is characterized by both stochastic coefficients and jumps. Transform methods are used to recover option prices. The model is calibrated to market data and the results compared to some well known stochastic volatility models.

1 Introduction

The Black-Scholes model for a stock is so commonly used as a starting point for analyses of derivative pricing, or optimal investment, that it is easy to miss the

*Wilberforce Road, Cambridge CB3 0WB, UK (phone = +44 1223 9798, e-mail = gd259@cam.ac.uk).

†Wilberforce Road, Cambridge CB3 0WB, UK (phone = +44 1223 766806, fax = +44 1223 337956, e-mail = L.C.G.Rogers@statslab.cam.ac.uk).

point that the price of an equity is a *derived* object, not a fundamental. Indeed, arbitrage-pricing theory tells us that the price of a stock is the net present value (NPV) of all future dividends, so once the dividend process has been specified (in the appropriate pricing measure) the price of the stock then follows. Though the Black-Scholes model is not normally presented in terms of a dividend process, it certainly can be (as we shall explain later), and this way of looking at the model can be extremely effective. As an example, we would cite the recent work of Korn & Rogers [KR] on options on stocks paying discrete dividends. By directly modelling the (discrete) dividend process of the stock, this paper is able to come up with a simple way of treating discrete dividends, free from inconsistencies, where various ‘industry’ modifications of the basic Black-Scholes equation have failed.

As we shall see, the Black-Scholes model with a constant dividend rate arises if we model the dividend process as a geometric Brownian motion in the pricing measure; so if we take *this* model for the dividend process, we will come up with nothing new. What we shall do instead is to take for the dividend process a *Markov-modulated* geometric Brownian motion, and derive the resulting dynamics of the stock price from an equilibrium analysis for a CRRA representative agent. Of course, the Black-Scholes model is a special case.

Markov-modulated dynamics have become increasingly popular in financial modelling over the past few years. Driffill, Kenc and Sola [DKS] take a similar approach to ours by modeling the the dividend process as a regime switching geometric Brownian motion. They then evaluate the stock as the net present value of future dividends, but in contrast to what we do here, they compute the expectation in the original measure, and not in the risk-neutral measure arising from equilibrium. They use PDE techniques to derive the price of shares and perpetual American calls. The present work generalizes the contribution of [DKS] by allowing the Markov chain driving the dividend dynamics to take N possible states, and by performing an equilibrium pricing. We also show how to price standard calls and puts (non-perpetual) by integral transform techniques. A similar model is studied by Veronesi [Ver], who again takes a two-state chain, but supposes that the state of the chain is not observable and has to be filtered from the observations. For this to make sense, Veronesi assumes that only the drift of the dividend process changes when the Markov state changes; our analysis assumes that the volatility as well as the drift will change when the Markov state changes, and this allows us to determine from the price exactly which state the chain is in at any time.

Jobert-Rogers [JR] derive the price of American options when stock prices (not dividends) follow a N -state Markov modulated diffusion. Konikov and Madan [KM] model the stock as a two state regime switching VG process. The characteristic function of the stock price is derived explicitly. Call and Put prices

are derived using Fourier transforms. Elliot and Osakwe [EO] extend the work of [KM] to a N-state Markov switching model where the underlying stock is a pure jump Levy process. Chourdakis [C2] assumes stock prices follow a Markov modulated Levy process (with continuous component) and explicitly derives the characteristic function of the share price. Vanilla option prices are also derived using FFT techniques. Our approach is different from the one taken in the above papers in that it considers the stock price as a derived quantity (via an equilibrium approach). In principle, one could assume some Levy dynamics for the underlying dividend and obtain a stock price processes similar to the one derived by [EO] and [C2]. However, our calibration shows that a simple Markov modulated GBM is already sufficient to explain the shape of the volatility surface. Finally Di Graziano and Rogers [DGR] extend the work presented in this paper to the pricing of various barrier style options.

In Section 2 we present the basic equilibrium model, and use it to derive the state-price density process, and hence the price of the stock. We shall suppose that the dividend process is modulated by a finite-state irreducible Markov chain; an interesting feature is that the price of the share jumps every time the underlying Markov chain jumps, in a way that is endogenously determined, rather than exogenously imposed. We also show how the Black-Scholes model results in the special case where the Markov chain takes only one value. It is worth commenting that even though the dividend process is continuous, the change of underlying Markov state results in jumps in the asset price, in contrast to models such as those of [JR] where it is the asset price itself that is modelled as a Markov-modulated process.

Next we go on in Section 3 to compute the prices of European put options; in contrast to the benchmark Black-Scholes model, there is no closed-form expression, and we have to resort to numerical methods (transform inversion) to evaluate option prices. This uses the method of Hosono [Ho], and Abate & Whitt [AW] to invert Laplace transforms.

The next stage is to calibrate the model (assuming a 2-state chain) to put option prices on the S&P 500 index. Fitting to three expiries and a range of moneyness, we find a calibrated model that fits data substantially better than various other non-Black-Scholes models in the literature¹, and shows remarkable intertemporal stability.

¹There is no standard calibration dataset available, so the other models have been calibrated to different datasets from ours.

2 The equilibrium model

We consider a very simple economy with a single productive asset, a stock, whose output (dividend) process δ obeys

$$\frac{d\delta_t}{\delta_t} = \mu(\xi_t)dt + \sigma(\xi_t)dW_t. \quad (1)$$

Here, μ and σ are deterministic functions of the irreducible finite-state Markov chain ξ , with Q -matrix Q , and W is a one dimensional standard Brownian motion independent of the Markov chain ξ .

In this economy, there is a single (representative) agent, who derives utility from consumption, and who owns the stock at time 0. This agent wishes to maximise his objective

$$E \left[\int_0^\infty e^{-\rho t} U(c_t) dt \right],$$

where the consumption process c must be chosen to satisfy the budget constraint

$$\zeta_0^{-1} E \left[\int_0^\infty \zeta_t c_t dt \right] = \zeta_0^{-1} E \left[\int_0^\infty \zeta_t \delta_t dt \right], \quad (2)$$

where ζ is the state-price density process. The left-hand side of (2) is the time-0 value of the agent's consumption process, whereas the right-hand side is the time-0 value of all the future dividends of the stock, that is to say, the price S_0 at time 0 of the stock. More generally, we have the price at time t of the stock is expressed as

$$S_t = \zeta_t^{-1} E_t \left[\int_t^\infty \zeta_u \delta_u du \right]. \quad (3)$$

It is well known (see, for example, Breeden [B], or Karatzas & Shreve [KS]) that the optimal consumption c^* is related to the state-price density process by

$$e^{-\rho t} U'(c_t^*) = \lambda \zeta_t \quad (4)$$

for some positive constant λ . Thus if markets are to clear, that is, $c^* = \delta$, we may use (4) to deduce what the state-price density process must be, up to an irrelevant positive constant.

For simplicity and tractability, we shall suppose that the agent has a CRRA utility

$$U(x) = \frac{x^{1-R}}{1-R},$$

for some positive R different from 1. In this case, we can express ζ explicitly in terms of δ , using (4) and the market-clearing condition:

$$\zeta_t = \lambda^{-1} e^{-\rho t} \delta_t^{-R}. \quad (5)$$

We may (and shall) without loss of generality suppose that $\lambda = 1$ in all that follows.

Within this very concrete model, it is possible to make the expression (3) for the stock price much more explicit. Indeed, taking $t = 0$ (with no real loss of generality) and using the notation $\tilde{\mu}(x) \equiv \mu(x) - \frac{1}{2}\sigma(x)^2$, we have

$$\begin{aligned}
S_0 &= \zeta_0^{-1} E_0 \left[\int_0^\infty \zeta_t \delta_t dt \right] \\
&= \delta_0^R E_0 \left[\int_0^\infty e^{-\rho t} \delta_t^{1-R} dt \right] \\
&= \delta_0 E_0 \left[\int_0^\infty \exp\left\{-\rho t + (1-R) \int_0^t \sigma(\xi_s) dW_s + (1-R) \int_0^t \tilde{\mu}(\xi_s) ds\right\} dt \right] \\
&= \delta_0 E_0 \left[\int_0^\infty \exp\left\{-\rho t + \int_0^t f(\xi_s) ds\right\} dt \right] \\
&\equiv \delta_0 v(\xi_0),
\end{aligned}$$

say, where

$$f(x) \equiv (1-R)\tilde{\mu}(x) + \frac{1}{2}(1-R)^2\sigma^2(x). \quad (6)$$

Thus we shall have in general that

$$S_t = \delta_t v(\xi_t), \quad (7)$$

where

$$v(x) = E \left[\int_0^\infty \exp\left\{-\rho t + \int_0^t f(\xi_s) ds\right\} dt \middle| \xi_0 = x \right].$$

Routine methods allow us to express v in terms of the generator Q of the Markov chain:

$$v = (\rho - Q - F)^{-1} \mathbf{1}, \quad (8)$$

where $\mathbf{1}$ is the constant vector all of whose entries are 1, and F is the diagonal matrix of f . We assume that ρ is large enough that all eigenvalues of $\rho - Q - F$ are in the open right half plane.

The explicit form (7) for the stock price shows how jumps of S arise at jump times of the chain ξ . The equilibrium analysis that has led to the equilibrium price (7) for S also allows us to compute the equilibrium riskless rate, and the martingale measure. Indeed, we have

$$\begin{aligned}
\zeta_t &= e^{-\rho t} \delta_t^{-R} \\
&= \exp\left\{-\rho t - R \int_0^t \sigma(\xi_s) dW_s - R \int_0^t \tilde{\mu}(\xi_s) ds\right\} \\
&= \exp\left\{-R \int_0^t \sigma(\xi_s) dW_s - \frac{1}{2} \int_0^t R^2 \sigma(\xi_s)^2 ds - \int_0^t r(\xi_s) ds\right\},
\end{aligned}$$

where

$$r(x) = \rho + R\tilde{\mu}(x) - \frac{1}{2}R^2\sigma^2(x). \quad (9)$$

Under the pricing measure we therefore have

$$\frac{d\delta_t}{\delta_t} = \sigma(\xi_t)dW_t^* + (\mu(\xi_t) - R\sigma(\xi_t)^2)dt, \quad (10)$$

where W^* is a Brownian motion in the pricing measure.

It is also clear from (7) that if $|I| = 1$ then the stock price process will be a geometric Brownian motion with no jumps, and with a dividend process proportional to S :

$$dS = S(\bar{\sigma}dW^* + (r - b)dt) \quad (11)$$

- that is, the Black-Scholes model. To match up the parameters in this familiar specification with those of the equilibrium model, we see from (7) and (10) that

$$\begin{aligned} \bar{\sigma} &= \sigma, \\ b &= \rho - f && \text{from (8), (7);} \\ &\equiv \rho - (1 - R)(\mu - R\sigma^2/2), \\ r &= \rho + \mu R - R(1 + R)\sigma^2/2 && \text{from (9),} \end{aligned}$$

and the consistency condition $r - b = \mu - R\sigma^2$ from comparing drifts in (10) and (11) is easily checked. There are three parameters in the equilibrium model, ρ , R and μ to fit the two parameters r , b of the Black-Scholes model, so there is one degree of indeterminacy; however, if we also require to match μ , then there is a unique choice of R and ρ that will work:

$$R = -\frac{r - b - \mu}{\sigma^2}, \quad (12)$$

$$\rho = b + (1 - R)(\mu - R\sigma^2/2). \quad (13)$$

There is no guarantee that these values of ρ and R will satisfy the required bounds $\rho > 0$ and $R > 0$, however. There is no contradiction here; an equilibrium model is arbitrage-free, but not every arbitrage-free pricing system need arise from an equilibrium, as this example shows.

3 Option pricing by integral transforms

In this section we show how to use integral transforms to price put options under the model described in the previous sections. It is well known (see, for example,

[CM], or [RZ]) that pricing a put option when we only know the cumulant-generating function of the log-stock price² is best done by transforming the put price in the log strike, and then numerically inverting the transform.

In more detail, the price of a put option with maturity T is given by

$$P_T(k) := \frac{1}{\zeta_0} E \left[\zeta_T (e^k - e^s)^+ \right], \quad (14)$$

where ζ_t is given by (5), and s and k are the log stock and the log strike respectively. For $\text{Re}(\alpha) > 1$ we define

$$\widehat{P}_T(\alpha) \equiv \int_{-\infty}^{\infty} e^{-\alpha k} P_T(k) dk, \quad (15)$$

and now change the order of integration, taking $\zeta_0 = 1$ with no loss of generality:

$$\begin{aligned} \widehat{P}_T(\alpha) &= \int_{-\infty}^{\infty} e^{-\alpha k} E \left[\zeta_T (e^k - e^s)^+ \right] dk \\ &= E \zeta_T \int_{-\infty}^{\infty} e^{-\alpha k} \int_s^{\infty} I_{\{y \leq k\}} e^y dy dk \\ &= E \zeta_T \int_s^{\infty} e^{-(\alpha-1)y} dy / \alpha \\ &= E \zeta_T \exp(-(\alpha-1)s) / \alpha(\alpha-1) \\ &= E[\zeta_T S_T^{1-\alpha}] / \alpha(\alpha-1) \\ &= \delta_0^{1-\alpha} e^{-\rho T} E \left[\delta_T^{1-\alpha-R} v(\xi_T)^{1-\alpha} | \delta_0 = 1 \right] / \alpha(\alpha-1) \\ &= \delta_0^{1-\alpha} e^{-\rho T} E \left[e^{(1-\alpha-R)(\int_0^T \sigma(\xi_s) dW_s + \int_0^T \tilde{\mu}(\xi_s) ds)} v(\xi_T)^{1-\alpha} | \delta_0 = 1 \right] / \alpha(\alpha-1) \\ &= \delta_0^{1-\alpha} e^{-\rho T} E \left[e^{\int_0^T z_\alpha(\xi_s) ds} v(\xi_T)^{1-\alpha} | \delta_0 = 1 \right] / \alpha(\alpha-1) \\ &= \delta_0^{1-\alpha} e^{-\rho T} \exp(T(Q + Z_\alpha)) v^{1-\alpha} / \alpha(\alpha-1), \end{aligned}$$

where

$$z_\alpha(x) \equiv (1 - \alpha - R)\tilde{\mu}(x) + \frac{1}{2}(1 - \alpha - R)^2 \sigma^2(x),$$

and Z_α is the diagonal matrix of z_α .

We used the method of Hosono [Ho] (see also Abate-Whitt [AW]) to numerically invert the Laplace transform of the put price (15). We then compared the results obtain with the B&S formula in case of constant dividend yield. Results are quite encouraging. The speed of computation using the Laplace transform method is broadly comparable with that of the analytic B&S formula.

²As in this example, or more commonly for log-Lévy price processes.

4 Hedging

The market we have so far analyzed is intrinsically incomplete, that is, apart from the trivial case of a single state chain, it is not possible to replicate all contingent claims in the economy (not even the simple put option considered in the previous paragraphs) by simply forming a self financing portfolio of stocks and bonds. In such a market in fact, we have no way to hedge the jump risk associated with a change in the state of the chain. However, we can *complete* the market by adding a number of traded options.

Assume that the stock price process is given by equation (3), where ξ_t is an N -state Markov chain. Suppose we want to hedge claim $C(t, \delta_t, \xi_t)$ and we have at our disposal instruments $H^{(1)}(t, \delta_t, \xi_t), \dots, H^{(N)}(t, \delta_t, \xi_t)$. In order to hedge the claim C , we form a self-financing portfolio V_t

$$V_t = \psi(t)B_t + \varphi^{(0)}(t)S(t) + \sum_{k=1}^N \varphi^{(k)}(t)H^{(k)}(t), \quad (16)$$

where $\varphi^{(k)}(t)$ is the quantity of asset k , with $k = 0, \dots, N$ held in the portfolio at time t , and $\psi(t)$ is the number of units of the money-market account B held at time t . We choose the portfolio weights $\varphi^{(i)}$ so that

$$C(t, \delta_t, \xi_t) = V_t \quad ;$$

by Itô's formula, this amounts to the conditions

$$\frac{\partial C}{\partial \delta} = \varphi^{(0)}(t)v(\xi_t) + \sum_{k=1}^N \varphi^{(k)}(t) \frac{\partial H^{(k)}}{\partial \delta} \quad (17)$$

$$\Delta C_{ij}(t) = \varphi^{(0)}(t)\Delta S_{ij}(t) + \sum_{k=1}^N \varphi^{(k)}(t)\Delta H_{ij}(t) \quad (18)$$

for $j = 1, \dots, N$. Here, we have written $\Delta C_{ij}(t) \equiv C(t, \delta_t, j) - C(t, \delta_t, i)$, with analogously $\Delta S_{ij}(t) = (v_j - v_i)\delta_t$. Equations (17) and (18) represent a system of $N + 1$ linear equations in the $N + 1$ unknowns $\varphi^{(i)}(t)$, $i = 1, \dots, N$, which can in general be solved to find the hedge ratios.

We still need to calculate the derivatives of the hedging instruments and claim C with respect to δ , and this requires knowledge of the structure of the claims in question. For example, in the case of a put, the analysis of the previous Section shows easily that

$$\begin{aligned} \bar{P}_T(\alpha) &\equiv \int_{-\infty}^{\infty} e^{-\alpha k} \frac{\partial P_T(\delta, k)}{\partial \delta} dk \\ &= -\frac{\delta_0^{-\alpha}}{\alpha} e^{-\rho T} \exp(T(Q + Z_\alpha)) v^{1-\alpha}, \end{aligned}$$

which is then inverted the usual way.

5 Calibration methodology

In order to understand how well the model we are working with might fit data, we took some put option prices on the S&P500, and tried a number of different fitting procedures. This section presents the calibration methods used, and the next presents and discusses the numerical results obtained.

The data set consisted of put option prices written on the S&P500 referring to five consecutive trading days from February 23th, 2004 to February 27th, 2004. The choice of this particular index was motivated by its liquidity and the availability of option data for a reasonable set of expiries and strikes.

For each day under consideration, the set of data consisted of 12 strikes ranging from 950 to 1200 (on average approximately -17% to +6% of the at the money strike), and 3 expiries: 18 September 04, 18 December 04 and 18 June 05. We shall indicate by \widehat{P}^n the set of data observed on day n . For each n thus, \widehat{P}^n is a 12×3 matrix whose ij^{th} entry is the price of put option with strike K_i and expiry M_j . We will denote by $P(\theta, x)$, the 12×3 matrix of model prices for parameter vector θ , given that the Markov chain is in state x . Explicitly, θ contains the coefficient of relative risk aversion R , the agent's discount factor ρ , the elements of the Q -matrix Q , and the vectors μ and σ from the model specification (1).

As a first attempt, we tried a day by day calibration using a two state Markov chain. In order to find the set of parameters θ_n which best fitted market data on day n , we minimized the Average Relative Percentage Error (ARPE)³

$$L(\theta_n) = \frac{1}{N} \sum_{ij} \frac{\left| \widehat{P}_{ij}^n - \left(\sum_{x=1}^S \pi_x^n P(\theta_n, x) \right)_{ij} \right|}{\widehat{P}_{ij}^n}, \quad (19)$$

where N is the total number of market prices considered in each day of the calibration, S the number of possible states of the Markov chain and π_x^n denotes the weight assigned to state x of the chain on day n . The ARPE criterion gives the average loss per dollar invested if we used the model prices instead of the market prices.

³The calibration was performed using also other quality of fit measures such as the Average Absolute Error as a percentage of the mean price (APE), the Average Absolute Error (AAE) and the Root Mean Square Error (RMSE). A mathematical definition of these measure is given in [SST]. The results obtained minimizing these loss functions are comparable to those obtained using ARPE, so we do not report them.

As it is apparent from (19), we are comparing observed prices with a weighted average of the prices stemming from our model. The rationale behind this is the following: the model requires knowledge of the current state of the Markov chain in order to produce a unique price for the option, and this state will in general not be observable. We propose therefore that the market price should be some average over the ‘pure’ prices obtained if the state of the Markov chain were known with certainty. Indeed, Arbitrage Pricing Theory tells us that market prices should be expectations (under the pricing measure) of discounted payoffs, and these expectations may be computed by firstly conditioning on the current state of the Markov chain, then averaging over the possible states of the chain; this explains the assumed form.

In calibrating the model thus, we have to mix over all possible states of the chain on each given day. To select the weights used, we should ideally use the conditional distribution (under the pricing measure) of the current state of the chain given all observations to date. In the absence of a long run of data to give a meaningful estimate of this distribution, we simply treat the weights as unknowns to be optimised over. This approach turned out to be the most effective and produced quite satisfactory results.

One of the main objections to the day-by-day calibration method is that the estimated parameters θ_n depend on n , so are not constant over the calibration period. However, as Table 1 shows, parameters estimated using this method were fairly stable over the (short) time interval. To increase parameter stability, we modified (19) by including a penalty for inter-day changes in the parameters. More precisely, we set

$$LS(\theta_n) = L^n(\theta_n) + \beta \| (\theta_n - \theta_{n-1})^2 \| . \quad (20)$$

where β is an arbitrary positive constant. A high β , will increase the stability of the parameter set but will worsen the fit to market data. In general we found that the stability of θ_n did not significantly improve by using (20) instead of (19), and the related results will be omitted.

So far we have only included option prices in the calibration, ignoring interest rates or equivalently bond prices. However, once we specify the parameter vector θ , we have uniquely determined the vector of interest rates consistent with the model. In other words, there exists only one piecewise constant, right-continuous, short-rate process $r(\xi)$ consistent with the model assumptions and it is given by

$$r(\xi) = \rho + \mu(\xi)R - \frac{1}{2}\sigma(\xi)^2R(R + 1); \quad (21)$$

see (9). Once the value of θ is chosen, the bond prices are given by

$$\begin{aligned} B^{t,T}(\theta, x) &= E_t \left[\exp \left(\int_t^T r(\xi_u) du \right) \right] \\ &= (\exp [(T-t)(Q - \mathbf{R})] \mathbf{1})(x), \end{aligned}$$

where \mathbf{R} is obtained by diagonalizing the vector r . Function (19) can be easily modified to include bond prices. The new quality of fit measure will be given by

$$LB(\theta_n) = \frac{L(\theta_n)}{2} + \frac{\gamma}{2N_B} \sum_T \frac{|\widehat{B}^{n,T} - \sum_{x=1}^N \pi_x^n B^{n,T}(\theta_n, x)|}{\widehat{B}^{n,T}}, \quad (22)$$

where γ is an arbitrary constant measuring the relative weight of bond prices in the calibration, $\widehat{B}^{n,T}$ is the market price of a US T-bill or T-bond on day n with maturity T , and N_B is the number of bond prices used in the calibration. In particular we tried to fit the short-term end of the US treasury curve by using the 3, 6, 12 month T-bill and the 2 year T-bond.

We also calibrated the model using a 3 state Markov-chain. However, the quality of fit was not better than the one obtained using a 2 state chain, and the estimation procedure was significantly slower. We thus omit these results.

6 Numerical Results

We shall now present some of the numerical results obtained using the calibration methodologies discussed in the previous section. Table 1 shows the value of the calibrated parameter set θ when π is chosen “optimally”, that is, it is included in the parameter set and no bond is included in the calibration.

Note that the average relative percentage price error is roughly of the order of 0.53%. In words, the model is able to fit the volatility surface skew quite closely. As a check we compared the implied volatility surface obtained from market prices with the one obtained using our model prices. To measure the goodness of fit of the surface, we used the following function

$$L_\sigma = \frac{1}{N} \sum_{ij} |\widehat{\sigma}_{ij}^n - \sigma_{ij}^n|.$$

The error so calculated was in the order of 6 bps, which is well inside the bid-ask spread for vanilla options on the S&P500. Also, estimated parameters appear to be fairly stable over time, even when we do not include any penalty for deviating from the previous day estimate. The numerical values obtained are broadly where

Table 1: Parameter estimates

	Day 1	Day 2	Day 3	Day 4	Day 5
μ_1	0.0460	0.0586	0.0585	0.0590	0.0562
μ_2	0.3397	0.3281	0.3278	0.3343	0.4056
σ_1	0.0859	0.1008	0.1019	0.1045	0.1075
σ_2	0.4423	0.4570	0.4557	0.4608	0.4867
Q_{12}	1.3702	1.0471	1.0291	1.0218	0.9425
Q_{21}	3.5778	3.5590	3.5608	3.5589	3.5638
ρ	0.1021	0.0900	0.0918	0.0911	0.0900
R	2.0072	2.1156	2.1067	2.1017	2.0648
π_1	0.0839	0.0849	0.0718	0.0619	0.0184
v_1	3.7133	4.0523	4.0024	4.0247	4.0036
v_2	3.6363	3.9955	3.9446	3.9647	3.8971
<i>ARPE</i>	0.626%	0.482%	0.556%	0.499%	0.466%

one would expect them to be, so the fit does not stretch belief. In particular, the estimates of v_1 and v_2 imply jumps in price of the order of 2% when the underlying state changes, and this is big enough to be noteworthy, but not so big as to be extremely rare.

Table 2 shows the parameter estimates when bonds are included in the calibration for both the approaches highlighted in the previous paragraph. As one would expect, trying to fit the interest rate curve and the volatility surface simultaneously leads to larger errors than in the ‘vanilla’ cases. However, the quality of fit is still acceptable.

Table 2: Parameter estimates including bonds

	Day 1	Day 2	Day 3	Day 4	Day 5
μ_1	0.0100	0.0100	0.0101	0.0100	0.0100
μ_2	0.0100	0.0106	0.0100	0.0143	0.0100
σ_1	0.0274	0.0299	0.0340	0.0345	0.0337
σ_2	0.2655	0.2572	0.2470	0.2482	0.2408
Q_{12}	0.3916	0.3603	0.3469	0.3466	0.3712
Q_{21}	3.4136	2.8707	2.6272	2.5873	2.5920
ρ	0.0900	0.0900	0.0900	0.0900	0.0900
R	4.3911	4.3038	4.3539	4.3397	4.3504
π_1	0.8382	0.8323	0.8325	0.8271	0.8274
<i>ARPE</i>	1.45%	1.37%	1.38%	1.34%	1.46%

The quality of fit is summarised in Figure 1, expressed in terms of the implied volatilities. Notice that this was not the criterion used to achieve the fit; nevertheless, the quality of fit is quite satisfying.

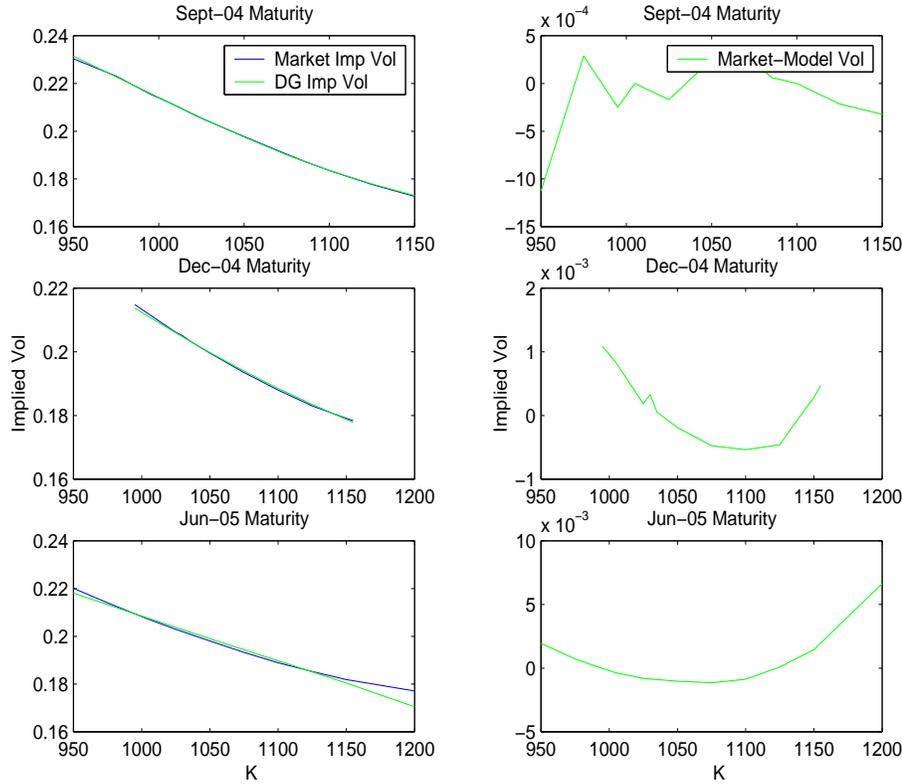


Figure 1: Summary of the quality of fit for each of the three expiries.

7 Comparison with other Models

In order to get a better grasp of how well our equilibrium model (for short, DR model), is able to fit market data and generate realistic dynamics for the volatility surface, we compared it to some well established equity models in the literature. In particular the benchmarks we chose were the Variance-Gamma and the Normal Inverse Gaussian model both under CIR and gamma-OU stochastic clock, which we shall refer to as VGCIR, VGOU, NIGCIR and NIGOU respectively (see Schoutens et al [SST]). Those models were calibrated on the same set of data described in section 5. For the purpose of this exercise, bond prices were not included in the data set.

As a first test, we performed a single-day calibration for all the models and compared the corresponding ARPE over the five day reference period. Table 3 shows that our DR model performs well, being the one associated with the lowest ARPE. It is conceivable that the better performance of the DR model can be partially attributed to its larger number of parameters compared to the other models (9 instead of 6).

Table 3: ARPE across models, single day calibration

Model	Day 1	Day 2	Day 3	Day 4	Day 5
VGOU	1.240%	0.979%	1.051%	1.026%	1.126%
VGCIR	0.922%	1.111%	1.049%	1.013%	1.095%
NIGOU	1.009%	0.954%	1.092%	1.004%	1.018%
NIGCIR	1.221%	0.919%	1.040%	0.977%	1.049%
DR	0.545%	0.462%	0.551%	0.501%	0.475%

While a higher number of parameters could be seen as a disadvantage when calibrating to a single volatility surface, the higher flexibility of the DR model (the number of parameters is a function of our choice of number of states in the Markov-chain) can prove extremely useful when calibrating to a larger data set, e.g. prices on i successive days. As a first step towards this direction, we fitted a single set of parameters to option price data relative to two consecutive trading days. More precisely, in the DR model with two states, we forced θ^n and θ^{n+1} to be the same while allowing π^n and π^{n+1} to be chosen optimally on each day. Similarly for the VGCIR, VGOU, NIGCIR and NIGOU we fixed the parameter set on any two consecutive days and allowed only the initial volatility v to vary. Results are reported in Table 4. Again, the DR model outperformed the others in terms of volatility surface fitting. In principle, the DR model allows you to calibrate to j days (say a week) simultaneously while keeping an acceptable fit by increasing the number of chain states, an option not present in the other models here examined. However, by increasing the number of chain states, the calibration effort becomes more involved and requires some extra care.

Table 4: ARPE across models, two days calibration

Model	Day 1	Day 2	Day 3	Day 4	Day 5
VGOU	1.217%	1.230%	1.789%	1.188%	2.112%
VGCIR	1.060%	0.989%	1.078%	1.139%	1.149%
NIGOU	1.078%	1.061%	1.122%	1.228%	1.177%
NIGCIR	1.134%	1.048%	2.427%	1.202%	1.183%
DR	0.690%	0.590%	0.593%	0.664%	0.582%

As a further check, we performed an out of sample test. More precisely, we used the set of parameters obtained from the day-to-day calibration on the n^{th} day to generate model prices for day $n + i$, where $i = 1 \dots 5$, and compared them to the corresponding market prices \hat{P}^{n+i} using (19). We repeated the same exercise in the case of the two-day fitting. Table 5 and Table 6 report the mean out of sample ARPE for each model. For example the 1-day frwd column represent the average error over the five day calibration period when we use today's parameters to compute tomorrow's prices.

One problem which arises when performing such an out of sample test, is how to choose the *forward* weights π^{n+i} and the *forward* initial volatility v^{n+i} relative to day $n+i$. For the sake of simplicity, in the case of the single day calibration we used π^n and v^n obtained from the n^{th} day calibration to generate the volatility surface for day $n+i$. In the case of the two days calibration, assuming that the calibration was performed on day $n-1$ and n , π^{n+1} and v^{n+i} were set equal to π^n and v^n respectively. As shown in Table 5 and Table 6 the out of sample error was comparable across models.

Table 5: Out of sample average ARPE, single day calibration

Model	1 day frwd	2 day frwd	3 day frwd	4 day frwd	5 day frwd
VGOU	2.440%	2.880%	3.770%	4.720%	5.090%
VGCIR	2.450%	2.730%	3.640%	4.560%	4.990%
NIGOU	2.520%	2.650%	3.690%	4.490%	4.860%
NIGCIR	2.490%	2.920%	3.800%	4.830%	5.170%
DR	2.310%	2.590%	3.530%	4.400%	4.770%

Table 6: Out of sample average ARPE, 2 day calibration

Model	1 day frwd	2 day frwd	3 day frwd	4 day frwd	5 day frwd
VGOU	2.540%	2.660%	3.840%	3.910%	4.840%
VGCIR	2.310%	2.240%	3.510%	3.410%	4.580%
NIGOU	2.360%	2.210%	3.460%	3.410%	4.650%
NIGCIR	2.350%	2.290%	3.670%	3.510%	4.640%
DR	2.130%	2.330%	3.610%	3.620%	4.580%

8 Conclusions

We have presented a simple equilibrium model where the single production activity in the economy pays a continuous stochastic dividend whose dynamics follows a Markov-modulated geometric Brownian motion. Such a model may offer one explanation for the observed phenomenon of jumps in stock prices. The model also allows for stochastic drift and volatility in the dynamics of the stock. We showed how to price European-style vanilla options by Laplace transform methods. In terms of computational speed, this method turned out to be very fast. Calibration to market data required some care, but even using a two state Markov chain, the model is able to fit the volatility surface quite closely.

References

- [AW] J. Abate and W. Whitt: Numerical inversion of Laplace transforms of probability distributions, *ORSA Journal on Computing* **7**, 36–43, 1995.
- [B] D. T. Breeden: Consumption, production, inflation and interest rates: a synthesis, *Journal of Financial Economics* **16**, 3–39, 1986.
- [CGMY] P. Carr, H. Geman, D. Madan, M. Yor: Stochastic Volatility for Levy Processes, *Mathematical Finance*, July 2003, Vol. 13, No. 3, 345-382.
- [CM] P. Carr and D. Madan: Option pricing and the fast Fourier transform, *Journal of Computational Finance* **2**, 61-73, 1999.
- [C] K. Chourdakis: Continuous-time regime switching models and application in estimating processes with stochastic volatility and jumps, Working paper, University of London, 2002.
- [C2] K. Chourdakis: Switching Levy Models in Continuous Time: Finite Distributions and Option Pricing, *SSRN*, working paper, 2006.
- [DGR] G. Di Graziano, L.C.G. Rogers: Barrier options for assets with Markov modulated dividends, *Journal of Computational Finance*, **9 – 4** 2006.
- [DS] J. Driffill and M. Sola: Option Pricing and Regime Switching, Birkbeck College, Department of Economics, 2000.
- [DKS] J. Driffill, T. Kenc, M. Sola: Merton-style option pricing under regime switching, *Computing in Economics and Finance*, 2002.
- [DBJ] G. Daniel, D. S. Bree, N.L. Joseph: Goodness of fit of the Heston model, *Computing in Economics and Finance Conference*, Society for Computational Economics, August 2003.
- [EO] R. Elliott and C. Osakwe: Option Pricing for Pure Jump Processes with Markov Switching Compensators, *Finance and Stochastics* **10**, 2006.
- [H] S. L. Heston: A closed form solution for options with stochastic volatility with applications to bond and currency options *Review of Financial Studies* **26**, 637-653, 2001.
- [Ho] T. Hosono: Numerical inversion of Laplace transform and some applications to wave optics, *Radio Science* **16**, 1015–1019, 1984.

- [JR] A. Jobert, L.C.G. Rogers: Option pricing with Markov-modulated dynamics, *SIAM Journal on Control and Optimization* **44**, 2063-2078, 2006.
- [KM] M. Konikov, D. Madan: Option Pricing Using Variance Gamma Markov Chains, *Review of Derivative Research*, **5**, 2001.
- [KS] I. Karatzas and S. E. Shreve: *Methods of Mathematical Finance*, Springer-Verlag, New York, 1998.
- [KR] R. Korn and L. C. G. Rogers: Stocks paying discrete dividends: modelling and option pricing. *The Journal of Derivatives* **13**, 44-48, 2005.
- [L] R. W. Lee: Option pricing by transform methods: extensions, unification and error control, *Journal of Computational Finance* **7**, 51-86, 2004.
- [P] P. Protter: *Stochastic Integration and Differential Equations*, Springer-Verlag, New York, 1990.
- [RW] L.C.G. Rogers and D. Williams: *Diffusions, Markov Processes and Martingales: Volume 2 , Ito's Calculus*, Cambridge University Press, Cambridge, 2000.
- [RZ] L.C.G. Rogers and O. Zane: Saddle-point approximations to option prices, *Annals of Applied Probability* **9**, 493–503, 1999.
- [SST] W. Schoutens, E. Simons, J. Tistaert: A perfect calibration! Now what?, Technical Report, K. U. Leuven, 2003.
- [Ver] P. Veronesi: Stock market overreaction to bad news in good times: a rational expectations equilibrium model, *Review of Financial Studies* **12**, 975–1007, 1999.