Implied volatility phenomena as market's aversion to risk

S.Siyanko, University College London, Gower Street, London WC1E 6BT, UK,
M. Pontier, Institut Mathématiques de Toulouse, France

June 27, 2017

Abstract

In this article, we show how to price derivatives in the presence of jumps and, in particular, when jumps are directed towards, what market participants believe, is the true, intrinsic value of an asset defined by its company fundamentals. Geometric Brownian motion is considered as a model for the asset price dynamics and jumps are introduced towards a deterministic exponential function. The function represents the fundamental value, and jumps are here to model price corrections in response to a new information available to the market. The objective is not only to introduce an alternative pricing formula but, more importantly, to have an algorithm that explicitly prices and measures the risk of price corrections assumed by the market.

Keyword: derivatives pricing, implied volatility, mean reversion through jumps, risk of price corrections.

1 Introduction

The Black-Scholes formula is one of the most renowned expressions in financial mathematics [2]. It is used by practitioners to calculate prices of standard call or put options which are respectively contracts to buy or sell an asset at a predefined price and at a predefined time. The formula requires six inputs - volatility of the underlying asset, its spot price, time to maturity, strike, risk-free interest rate, and dividend yield.

At first sight, it may be somewhat of a surprise that option prices only depend on the spot price but not on any other information about the underlying such as its expected return. After all, one might argue that such information should clearly influence buy decisions because nobody would wish to purchase an option that is likely to be out of the money at its maturity. This in turn should influence the demand but somehow it does not happen in the Black-Scholes analysis. Of course, those adept in financial mathematics would note that it is not the expected value of the underlying but rather the possibility to replicate the option's final payoff that sets the price. It happens that one can sell an option and then meet one's obligations at maturity by actively managing a self-financing portfolio constructed from the underlying asset and a money market account. It is then the initial cost of this portfolio that sets the price. Indeed, if this was not the case, it would be possible to create an arbitrage strategy by taking opposite positions in the replicating portfolio and the option (you would need to short the option if it is overpriced or buy it if it is underpriced). It is then the willingness of large institutions to do arbitrage that will fuel demand and supply in the direction of the Black-Scholes formula and make the option price indifferent to the expected rate of return on the underlying asset.

One of the key assumptions behind the Black-Scholes model is that perfect replication is possible or, at least, we can replicate the payoff "on average". This is usually referred to as market completeness. In reality, the markets are not complete, and spot prices are not the only information about the underlying that markets respond to. Expected return on an asset can be estimated by market analysts through so-called fundamental analysis which takes into account the firm's income generating ability. This clearly is of high importance for hedge funds and indeed investors in general. At the same time, we know that there is technical trading going on which can make assets over or undervalued with respect to their true, intrinsic value defined by the fundamental analysis. When this happens, market participants may expect a correction in the asset price, and the correction would be sudden and unpredictable (otherwise, one could easily make money by buying an asset when it is undervalued and selling when it is overvalued). It is worth noting that the price correction may not happen at all (after all, fundamental analysis is not exactly a precise science) but it is the expectation of this correction that is important. Clearly, option sellers would see the possibility of price corrections as a risk. For example, if market participants expect an asset price to suddenly go up, the writer of a call option will be cautious about selling it cheap because he may not be able to hedge the risk of facing a higher payout by immediately readjusting positions in the underlying. This risk would obviously attract some premium, and, since there is a higher demand for this option (it is more likely to end up in the money) and limited arbitrage opportunities (the market is not complete), we have a new price that is higher than the price predicted by the Black-Scholes model.

We hypothesize that a premium for the risk of price corrections is indeed included in option prices, and it manifests itself through market phenomena called volatility skew and volatility smile. While, for the purpose of this research, this is an assumption, there is some evidence to suggest that the hypothesis is reasonable. Web sites targeted at investors often encourage them to interpret volatility skews as a sign of a "bullish" or "bearish" market in the underlying, and it is clear that both behaviours, being speculative in nature, indicate misalignment of the stock price with its fundamental value. The volatility smile phenomenon may also appear after a sharp price correction in the underlying (e.g. following a financial crisis) which supports the hypothesis because one would expect the volatility smile to appear when the stock price is close to what markets believe is its true, fundamental value.

There are, of course, many models that would have a good fit to volatility skews and smiles. Perhaps the most known and well used in the industry would be the Heston model [8]. It has two correlated Brownian motions driving the square of stochastic volatility and the price process itself. The stochastic volatility process is a mean-reverting process, and the motivation comes from the analysis of the statistics of log returns. It is known that the physical volatility of a stock normally exhibits volatility clustering (small moves are more likely to be followed by small moves and large moves by large) suggesting autocorrelation and thus motivating the mean reversion because mean-reverting processes would normally exhibit autocorrelation. In contrast, we will have a mean reversion that does not rely on or describe statistical properties of stock prices. It is fundamental that derivatives prices in our model will be formed not by a response to physical stock price movements but by a fear that sharp corrections in the stock price may occur. Having large jumps with low probability may simply not manifest itself in the statistics of log returns for years but it is clear that such events may have a significant influence over derivatives prices. It is also interesting that fitting volatility surface to the Heston model means that it does price risk but prices it exactly at the same level as everybody else. This is because, if model parameters are not independent of the model, the only way to estimate them is through calibration to historic derivatives prices. If we are to price risks differently from the market, the derivatives would either be too expensive to sell (if we are overpricing risks) or we would expose ourselves to greater risks than the other financial institutions. That would lead us to conclude that models that are pricing risks explicitly (or, to that matter, just exposing them) are not going to be very popular in the industry. They would simply not serve the business objective of pricing risks at the same level as everybody else does it. It is only when one wants to measure risks assumed by the market or/and, perhaps, to price risks in a more responsible way, our model would have its superior utility.

Mean reversion through jumps is used to model sharp corrections in the stock price, and there are many models that have this feature as well. This is particularity common when we consider assets that can not be stored, e.g. electricity. The non-storability of electricity means that instantaneous shocks to demand or supply (e.g. increase in demand due to a heat wave) are not smoothed by inventories and will result in shocks to the spot price followed by mean reversion to levels determined by the marginal costs of generating it. An example of such model would be [7]. Here, we have jumps and their direction (but not size) depends on upper and lower price thresholds in a way that leads to mean reversion.

There is also a research by [1] which are particularly useful. It happens that the problem for derivatives pricing was studied for a very wide class of diffusion processes with jumps, and the processes are similar to the risk-neutral stock price dynamics in our model. However, there are technical constraints that ensure that solutions are of an exponential type. The model here can be viewed as an extension to this framework leading to very different type of solutions - solutions of a non-exponential type.

Finally, there is a model that has been introduced by [13]. It is actually a special case of our model but in the context of pricing default risks - it is the Merton's jump-to-ruin model. If we set the fundamental value of the stock to zero, the risk of the price correction is effectively the same as the risk of default. While this is slightly out of context, [6] does comment on the suitability of this model to describe the left wing (low strike structure) volatility skew and, in particular, for the stocks with high credit spreads. He has also commented that the model does not generate the right wing which is just not consistent with the data. Since we will have both upside and downside risk of the price correction, our model will produce both wings of the volatility surface, and one can view it just as a natural extension to the Merton's jump-to-ruin model.

We start by introducing mean-reverting jumps into the geometric Brownian motion model (Section 2), and show that this price dynamics does not lead to a reasonable arbitrage-free market but motivates a different model. The new model is arbitrage-free and leads to a pricing PDE that we can use to price derivatives. Surprisingly, we obtain the same pricing PDE through the local no-arbitrage argument and with a very reasonable replicating strategy applied to the initial price dynamics (Section 3). Then, we simplify our PDE, derive its border conditions, present some properties of the solution and show how to solve it numerically (Section 4). Finally, we present numeric results (Section 5) followed by some practical advice for pricing applications (Section 6) and the conclusion (Section 7).

2 Stock price dynamics

In the Black-Scholes framework for pricing derivatives, we have standard geometric Brownian motion as a model for stock price dynamics. We can also use it but only to model stock price between price corrections, and, as for the price corrections, we can have jumps at random times defined by a standard Poisson process. The jumps should be directed towards what people think is the fundamental value of the stock, and we can model it as a standard exponential function. We can also consider the same model but with a compensated Poisson process. While it is not clear how it makes financial sense, compensated Poisson process is a martingale which is often desirable when pricing derivatives.

There is one property that we would like our models to have. It comes from the notion that price corrections are rear events that are managed and priced as risk. We want to ensure that we obtain a pricing algorithm that does not contradict risk management behaviors observed in the market. After all, we want an algorithm that explicitly prices and measures the risk of price corrections, not just another algorithm for pricing derivatives.

We are therefore going to introduce two models. One with a standard Poisson process and the other with its compensated version. We will start with the first model and show how, in light of the above requirement, it motivates the introduction of the second.

2.1 Geometric Brownian motion with Poisson jumps

Let us define a filtered probability space $\{\Omega, \mathcal{A}, \mathbb{F} = (\mathcal{F}_t, 0 \leq t \leq T), P\}$ on which we have an independent pair of a standard Brownian motion W and a Poisson process N with constant intensity λ. Here, we present our first attempt to model the stock price with mean-reverting jumps:

$$(1)\quad dS_t = \mu S_{t-}dt + \sigma S_{t-}dW_t + (\overline{S}_0 e^{\mu t} - S_{t-})dN_t,\ S_0 \in L^1(\Omega, \mathcal{F}_0, P),\ 0 \leq t \leq \text{ maturity } T,$$

where S_0, W, N are independent, and W, N are $\mathbb{F}$-adapted. Here, we are using standard notations for the stock price S_t, its drift rate μ, volatility σ, and time t. These are inputs for the geometric Brownian motion model. In addition, we have the standard Poisson process N as an extra source of randomness and the fundamental value of the stock at time t: $\overline{S}_0 e^{\mu t} = \overline{S}_t$ with constant $\overline{S}_0$. At jump times $(T_n, n = 1, 2, ...)$, the last term in the above equation introduces an increment equal to the difference between the stock price and its fundamental value, thus simulating price corrections.

Let us remark that this SDE is solvable and admits the unique positive solution:

$$(2)\quad \begin{aligned} S_t &= S_0 \exp\left[(\mu - \sigma^2/2)t + \sigma W_t\right] > 0, T_{N_t} = 0 \\ S_t &= \overline{S}_0 e^{\mu T_{N_t}} \exp\left[(\mu - \sigma^2/2)(t - T_{N_t}) + \sigma(W_t - W_{T_{N_t}})\right] > 0, T_{N_t} > 0, \end{aligned}$$

i.e. we have a standard geometric Brownian motion evolving from its initial value at time zero or from its fundamental value at jump time T_{N_t}.

Proposition 2.1 *There exist $\mathbb{F}$-adapted processes ψ, γ, $(\gamma > -1)$ which define a risk-neutral probability measure:*

$$Q = L_T P,\ \textit{where}\ L_t = \rho_t^W \rho_t^N,$$

where

$$d\rho_t^W = \rho_t^W \psi_t dW_t,\quad d\rho_t^N = \rho_{t-}^N \gamma_t (dN_t - \lambda dt).$$

So this market is viable, there is no arbitrage opportunity.

Proof: Let us now look at how we can price derivatives and, following the standard risk-neutral valuation framework, we start with a change of probability measure that will make our process "risk-neutral", i.e. there exist martingales $\tilde{W}$ and $\tilde{M}$ such that

$$(3)\quad dS_t = rS_{t-}dt + \sigma S_{t-}d\tilde{W}_t + (\overline{S}_0 e^{\mu t} - S_{t-})d\tilde{M}_t,\ S_0,\ t \leq \text{ maturity } T.$$

It means that we are looking for a pair ψ, γ such that the following processes are martingales:

$$
\begin{aligned}
d\rho_t^W &= \rho_t^W \psi_t dW_t \\
d\rho_t^N &= \rho_{t-}^N \gamma_t (dN_t - \lambda dt),
\end{aligned} \tag{4}
$$

and from Girsanov's Theorem (cf. Jeanblanc Yor page 69 and page 478 [10]) we know that if the process $L = \rho^W \rho^N$ is a martingale then under $Q = L.P$:

$$
\begin{aligned}
\tilde{W}_t &= W_t - \int_0^t \psi_s ds \\
\tilde{M}_t &= N_t - \int_0^t \lambda(1+\gamma_s) ds,
\end{aligned}
$$

where $\tilde{W}$ is a Q Brownian motion and N admits the intensity $\lambda(1+\gamma_t)$ under Q. Substituting W and N in (1) yields:

$$
dS_t = \mu S_{t-} dt + \sigma S_{t-}(d\tilde{W}_t + \psi_t dt) + (\overline{S}_0 e^{\mu t} - S_{t-})(d\tilde{M}_t + \lambda(1+\gamma_t)dt),
$$

and we deduce (3) if the following equation is satisfied:

$$
r = \mu + \sigma \psi_t + \frac{\overline{S}_0 e^{\mu t} - S_{t-}}{S_{t-}} \lambda(1+\gamma_t). \tag{5}
$$

Since we do not have a unique solution for ψ and γ (the market is not complete), let us propose:

$$
1 + \gamma_t = \inf(S_{t-}, c), \tag{6}
$$

for some constant $c > 0$. This leads us to the following expression for ψ:

$$
\psi_t = \frac{r-\mu}{\sigma} + \left(1 - \frac{\overline{S}_0 e^{\mu t}}{S_{t-}}\right) \frac{\lambda \inf(S_{t-}, c)}{\sigma}.
$$

Such a ψ is bounded on the interval $[0, T]$ because $0 \leq \frac{\inf(S_{t-}, c)}{S_{t-}} \leq 1$. From Theorem II.2.a in [11] we know that ρ^W is a square integrable martingale. Let us use M to denote the martingale $\int_0^t \gamma_s (dN_s - \lambda dt)$ and confirm from expression (6) that the predictable compensator of $\int_0^t \gamma_s dN_s$, which is $\lambda \int_0^t \gamma_s ds$, is bounded on $[0, T]$. It means (cf. Theorem II.2.b in [11]) that ρ^N is a martingale with integrable variation.

We therefore have an equivalent probability measure Q, under this measure the intensity of price corrections λ_t is given by

$$
\lambda_t = (1+\gamma_t)\lambda, \tag{7}
$$

where γ_t is defined in (6). •

It is interesting that λ_t is a constant for $S_{t-} \geq c$, and this is what people assume when pricing risks. But the drawback is that λ depends on S_{t-} when S_{t-} is below c. Does it make any financial sense, and if not why does it happen? We know from equation (2) that $S_t > 0$ P almost surely, so also Q almost surely because probability measures P and Q are equivalent.

Nevertheless, what actually stops us from postulating that, when pricing risk, one is using the physical probability measure with constant λ? It seems like a reasonable assumption to make. That would mean that $\gamma_t = 0$, and the associated process L could be only a local martingale, because we do not have easy sufficient condition on $\psi_t = \frac{r-\mu}{\sigma} + (1 - \frac{\overline{S}_0 e^{\mu t}}{S_{t-}})\frac{\lambda}{\sigma}$ (as Novikov's condition) to be satisfied. And if L is not a martingale, we can not define the equivalent probability measure Q, and (under this constraint) the market may allow an arbitrage which is fine. Technical trading does mean arbitrage, this is how technical traders make money, and the introduction of technical trading is indeed our objective here. On a positive side, it leads to a much simpler pricing rule defined by (3) but with constant jump intensity λ. Moreover, we can redefine our filtered probability space and say that this is our process under a measure Q (different from the Q above). Then we can introduce a change of probability measure between P and Q by stating that $\gamma_t = 0$, $\psi_t = (r-\mu)/\sigma$, and we are back to the standard framework of financial mathematics. The only problem is that our newly defined process allows negative values but there are many models that do this, we just need to postulate that our case of interest is when the probability of such events is low.

2.2 The Model

We are operating in a filtered probability space $\{\Omega, \mathcal{A}, (\mathcal{F}_t, 0 \le t \le T), P\}$ endowed with a filtration $\mathbb{F}$ generated by a standard Brownian motion W and a Poisson process N with constant intensity λ and independent of W. We can write our stock price dynamics under P as:

$$dS_t = \mu S_{t-}dt + \sigma S_{t-}d\tilde{W}_t + (\overline{S}_0 e^{\mu t} - S_{t-})(dN_t - \lambda dt),\ S_0 \in \mathbb{R}^+,\ 0 \le t \le \text{ maturity } T.$$

We now define an equivalent probability measures $Q = L_T P$ with the density of probability measure $dL_t = -L_t \frac{\mu - r}{\sigma} d\tilde{W}_t$ and under Q we have:

$$dS_t = rS_{t-}dt + \sigma S_{t-}dW_t + (\overline{S}_0 e^{\mu t} - S_{t-})(dN_t - \lambda dt), \tag{8}$$

where $dW_t = d\tilde{W}_t + \frac{\mu - r}{\sigma}dt$. Here we only change the Brownian motion part, and under both probability measures P and Q the intensity λ is the same. As in the previous model, there exists a risk-neutral probability measure, namely Q, so once again this market is viable, there is AOA (absence of opportunity arbitrage).

Proposition 2.2 *The unique solution to (8) under probability measure Q, before the first N jump time T_1 -meaning on the event $\{T_{N_t} = 0\}$, is*

$$S_t = e^{(r+\lambda-\frac{\sigma^2}{2})t+\sigma W_t}\left(S_0 - \lambda\int_0^t \overline{S}_0 e^{\mu s} e^{-(r+\lambda-\frac{\sigma^2}{2})s-\sigma W_s} ds\right) \tag{9}$$

and similarly, on the event $\{T_{N_t} > 0\}$, the solution is:

(10)

$$S_t = e^{(r+\lambda-\frac{\sigma^2}{2})(t-T_{N_t})+\sigma(W_t - W_{T_{N_t}})}\left(\overline{S}_0 e^{\mu T_{N_t}} - \lambda\int_{T_{N_t}}^t \overline{S}_0 e^{\mu s} e^{-(r+\lambda-\frac{\sigma^2}{2})(s-T_{N_t})-\sigma(W_s - W_{T_{N_t}})} ds\right).$$

Proof: We can check, via Itô's formula, that, before the first N jump time T_1, the unique solution to (8) under probability measure Q is

$$S_t = Y_t A_t,$$

where $A_t = S_0 - \lambda\int_0^t \overline{S}_0 e^{\mu s} Y_s^{-1} ds$ and $Y_t = \exp[(r + \lambda - \sigma^2/2)t + \sigma W_t]$.
Indeed, this is a product of Y_t satisfying

$$dY_t = Y_t[(r + \lambda)dt + \sigma dW_t]$$

and the finite variation process A satisfying $dA_t = -\lambda \overline{S}_0 e^{\mu t} Y_t^{-1} dt$. So the differential of the candidate solution is

$$Y_t dA_t + A_t dY_t = -Y_t \lambda \overline{S}_0 e^{\mu t} Y_t^{-1} dt + A_t Y_t[(r+\lambda)dt + \sigma dW_t] = A_t Y_t[(r+\lambda)dt + \sigma dW_t] - \lambda \overline{S}_0 e^{\mu t} dt.$$

Thus, the explicit solution before the first N jump time T_1 is

$$S_t = e^{(r+\lambda-\frac{\sigma^2}{2})t+\sigma W_t} \left(S_0 - \lambda \int_0^t \overline{S}_0 e^{\mu s} e^{-(r+\lambda-\frac{\sigma^2}{2})s - \sigma W_s} ds \right)$$

on the event $\{T_{N_t} = 0\}$, and similarly, on the event $\{T_{N_t} > 0\}$, we can show that the solution is:

$$S_t = e^{(r+\lambda-\frac{\sigma^2}{2})(t-T_{N_t})+\sigma(W_t - W_{T_{N_t}})} \left(\overline{S}_0 e^{\mu T_{N_t}} - \lambda \int_{T_{N_t}}^t \overline{S}_0 e^{\mu s} e^{-(r+\lambda-\frac{\sigma^2}{2})(s-T_{N_t}) - \sigma(W_s - W_{T_{N_t}})} ds \right).$$

•

Note that, since the process $\lambda \int_0^t \overline{S}_0 e^{\mu s} e^{-(r+\lambda-\frac{\sigma^2}{2})s-\sigma W_s} ds$ is continuous and strictly increasing, S_t may take negative values, e.g. after time

$$\tau := \inf\{t \geq 0 : \quad \lambda \int_0^t \overline{S}_0 e^{\mu s} e^{-(r+\lambda-\frac{\sigma^2}{2})s-\sigma W_s} ds \leq S_0\} \tag{11}$$

on the event $\{\tau < T_1\}$ and, more generally, between two jump times when

$$\lambda \int_{T_{N_t}}^t \overline{S}_0 e^{\mu s} e^{-(r+\lambda-\frac{\sigma^2}{2})(s-T_{N_t})-\sigma(W_s - W_{T_{N_t}})} ds \leq \overline{S}_0 e^{\mu T_{N_t}}.$$

This is not desirable when pricing derivatives but, on a positive side, we believe that having constant λ better reflects pricing behaviors in the market, and it also helps with tractability of the solution. Moreover, λ as a constant means that there is a very reasonable replicating strategy that we can apply to our initial price dynamics (1) and, with the help of the local no-arbitrage argument, derive the same pricing PDE as in this model.

That concludes the introduction of the model for the stock price dynamics. We can now proceed to derive the pricing PDE.

3 Deriving the pricing PDE

3.1 First approach

Throughout this section we are operating under the risk-neutral probability measure Q. Let us start by noting that the process (8) is a Markov process which leads us to conclude that the option's price C_t is a function of S_t.

Indeed, we can invoke the fundamental theorem of asset pricing and write:

$$C_t = e^{-r(T-t)} E_Q[(S_T - K)^+ / \mathcal{F}_t], \tag{12}$$

and, since S is a Markov process, there exists a regular function C such that

$$C_t = C(t, S_t). \tag{13}$$

Proposition 3.1 *Assuming the function C defined on $[0,T]\times\mathbb{R}$ is $C^{1,2}$, it is solution to the partial derivatives equation*

(14)

$$\frac{\partial C}{\partial t}(t,s)+\frac{1}{2}\sigma^2 s^2\frac{\partial^2 C}{\partial s^2}(t,s)=r[C(t,s)-s\frac{\partial C}{\partial s}(t,s)]-\lambda[(C(t,\overline{S}_0e^{\mu t})-C(t,s))-\frac{\partial C}{\partial s}(t,s)(\overline{S}_0e^{\mu t}-s)],$$

with boundary conditions

$$C(T,s)=(s-K)^+,\ \ \partial_s C(t,s)=0 \text{ if } s\leq 0. \tag{15}$$

Proof: The process $t\to e^{r(T-t)}C(t,S_t)$ is an $(\mathbb{F},Q)$-martingale and using Itô's formula we cancel its finite variation part:

$$\frac{\partial C}{\partial t}+rS_{t-}\frac{\partial C}{\partial s}+\frac{1}{2}\sigma^2S^2_{t-}\frac{\partial^2 C}{\partial s^2}+\lambda[(C(t,\overline{S}_0e^{\mu t})-C)-(\overline{S}_0e^{\mu t}-S_{t-})\frac{\partial C}{\partial s}]=rC(t,S_t)\ dt\otimes dQ\ a.s..$$

The stochastic process S admits the support $\mathbb{R}$, so we can replace $S_t(\omega)$ by $s\in\mathbb{R}$ and obtain the PDE on $\mathbb{R}^+\times\mathbb{R}$:

$$\frac{\partial C}{\partial t}(t,s)+\frac{1}{2}\sigma^2 s^2\frac{\partial^2 C}{\partial s^2}(t,s)=r[C(t,s)-s\frac{\partial C}{\partial s}(t,s)]-\lambda[(C(t,\overline{S}_0e^{\mu t})-C(t,s))-\frac{\partial C}{\partial s}(t,s)(\overline{S}_0e^{\mu t}-s)],$$

where we have $t\in[0,T)$, $s\in(-\infty,+\infty)$, and the terminal condition is:

$$C(T,s)=(s-K)^+.$$

•

While border conditions (at $s=0$ and $s\to+\infty$) are not necessary to define the PDE, they are useful for numeric methods, and we will derive them later.

3.2 Through the local no-arbitrage argument.

It is interesting that we can derive the same PDE from the initial price dynamics (1):

$$dS_t=\mu S_{t-}dt+\sigma S_{t-}dW_t+(\overline{S}_0e^{\mu t}-S_{t-})dN_t \tag{16}$$

Note that, formally speaking, this is not our process process under Q or P because (12) prices the option under Q, not P. This is the price dynamics we have started from in order to motivate the introduction of our model stated in the previous SDE (1).

Let us start with looking for a self-financing portfolio:

$$V_t=\phi_tS_t+\psi_tB_t \tag{17}$$

with some predictable processes ϕ and ψ denoting our positions in stock and the money market account. We now try to replicate the option over a small time interval dt. Recalling the assumption that the option price process C is a regular function of (t,S_t) and with the help of Itô's lemma we obtain:

$$dC(t,S_t)=\frac{\partial C}{\partial t}dt+\mu S_t\frac{\partial C}{\partial s}dt+S_t\frac{\partial C}{\partial s}\sigma dW_t+\frac{1}{2}\sigma^2S_t^2\frac{\partial^2 C}{\partial s^2}dt+[C(t,\overline{S_0}e^{\mu t})-C(t,S_{t-})]dN_t. \tag{18}$$

On the other hand, we know that for the self-financing portfolio V_t we can write:

$$dV_t=\phi_tdS_t+\psi_tdB_t=\phi_tdS_t+r(V_t-\phi_tS_t)dt. \tag{19}$$

Then, we assume that C and V admit the same finite variation part $dt \otimes dQ$ almost surely on $[0, T) \times \Omega$. Thus, we delta hedge: $\phi_t = \frac{\partial C}{\partial s}(t-, S_{t-})$ between jumps. Remark that with such a portfolio we equate Brownian motion coefficients on both sides. However, we can not equate coefficients of Poisson process: jumps will introduce a misalignment between the value of the replicating portfolio and the call option equal to:

$$C(t, \overline{S}_0 e^{\mu t}) - C(t-, S_{t-}) - \phi_t(\overline{S}_0 e^{\mu t} - S_{t-}) \tag{20}$$

meaning that jumps are managed as a risk, i.e. we are using the physical probability measure and replicate "on average":

$$E_t[dC_t - dV_t] = 0. \tag{21}$$

That is, if a jump happens, one needs to immediately re-adjust his positions in the stock price (as per the delta hedge) and the money market account. This means that we can not replicate the option's price process exactly. Best we can do is to ensure that we replicate the payoff "on average":

$$V_t = C(t, S_t) \tag{22}$$

$$E[C(T, S_T)/\mathcal{F}_t] = E[V_T/\mathcal{F}_t]. \tag{23}$$

Since the above would hold true for martingales, we just need to identify and equate finite variation parts of semi-martingales (18) and (19) which yields:

$$\begin{aligned} &\frac{\partial C}{\partial t} + \mu S_{t-}\frac{\partial C}{\partial s} + \frac{1}{2}\sigma^2 S_{t-}^2 \frac{\partial^2 C}{\partial s^2} + \lambda[C(t, \overline{S_t}) - C(t, S_{t-})] = \\ &\mu S_{t-}\frac{\partial C}{\partial s} + \lambda(\overline{S}_0 e^{\mu t} - S_{t-})\frac{\partial C}{\partial s} + r(C - S_{t-}\frac{\partial C}{\partial s}), \end{aligned} \tag{24}$$

and after cancellation we recover the PDE (14).

We have therefore presented a replicating strategy for our initial price dynamics (1) leading to the same PDE as in the no-arbitrage market model exhibited in Section (2.2). The strategy is very reasonable - markets do delta hedge and ignoring jumps means that we are over or under inflating (vs. Black-Scholes formula) our position in the money market account with a view to attain the payoff "on average". Moreover, from (20) we can see that, if $C(t, s)$ is a convex function of s, we are over inflating our position in the money market account to ensure that we can recover losses from jumps, i.e. we delta hedge and manage the possibility of jumps as a risk.

Remark 1 *Our initial price dynamics (1) is strictly positive but the PDE is defined on $s \in R$. This hints at a possibility to restrict it to $s \in R^+ - \{0\}$. Note, such restriction is not possible in general, and this result is not trivial.*

4 Pricing PDE

Let us now move on to present the pricing PDE and explore some of its properties. For simplicity, we limit the scope to standard vanilla call options and postulate that we operate in a risk-free environment with the pricing rule (12) and the stock price process (8). With the help of Itô's lemma, we have derived PDE (14), and we can now simplify it further by introducing an auxiliary function u on $[0,T]\times\mathbb{R}^+$ satisfying:

$$C(t,s)=u(t,s)e^{-(r+\lambda)(T-t)}+\lambda e^{-r(T-t)}\int_t^T e^{-\lambda(T-v)}u(v,\overline{S}_0e^{\mu v})dv. \tag{25}$$

Proposition 4.1 *Assuming the existence of a function u satisfying (25), the PDE (14) is equivalent to the following PDE with u being a solution of it:*

$$\frac{\partial u}{\partial t}+\frac{1}{2}\sigma^2s^2\frac{\partial^2 u}{\partial s^2}=-s(r+\lambda)\frac{\partial u}{\partial s}+\lambda\overline{S}_0e^{\mu t}\frac{\partial u}{\partial s},\quad u(T,s)=(s-K)^+. \tag{26}$$

Proof: We can use the expression for $C(t,s)$ (25) to calculate its derivatives $\frac{\partial C}{\partial t},\frac{\partial C}{\partial s},\frac{\partial^2 C}{\partial s^2}$ and $C(t,\overline{S}_0e^{\mu t})$ with respect to the derivatives of the fnction u. Then, plugging all of these derivatives into the PDE (14) and its terminal condition (15) leads to

$$\frac{\partial u}{\partial t}+\frac{1}{2}\sigma^2s^2\frac{\partial^2 u}{\partial s^2}=-s(r+\lambda)\frac{\partial u}{\partial s}+\lambda\overline{S}_0e^{\mu t}\frac{\partial u}{\partial s}$$

and obviously $u(T,s)=C(T,s)=(s-K)^+$.
Conversely, if u is a solution of (26), the function C defined by (25) satisfies the original PDE (24). •

When working with jumps, it is normal to end up with equations of an integro-differential type. The "integro" part comes from changes in the derivative's price following jumps, and it could significantly complicate the solution. It is interesting that there is no "integro" part in the above PDE (26) - we have removed it by substituting $C(t,s)$ with $u(t,s)$. This is actually a special case, it is not possible to remove the "integro" component in general.

As for the solution to the PDE itself, we have a standard Black-Scholes Heat equation but with a convection term arising from the drift $-\lambda(\overline{S}_0e^{\mu t}-S_t)dt$ in the price dynamics (8). The drift has a non-proportional component with respect to the stock price S_t which significantly undermines analytical tractability of the PDE solution.

A similar situation would arise, for example, if we have non-proportional dividends paid on a stock. While this is how dividends are normally paid, it is often assumed that they are proportional to the stock price, and analytical formulae can be derived and used as good enough estimates. Unfortunately, if we are to assume that the fundamental value of the stock is proportional to its spot price, we are going to remove the mean reversion.

If we are to solve the PDE (26) as it is, we need to think of a case where similar equations may appear. Using standard Feynman-Kac's formula, e.g. [10] 10.2.5, we get the corollary

Corollary 4.2 *The unique solution to (26) is the following*

$$u(t,s)=E[(X_T^{t,s}-K)^+], \tag{27}$$

where

$$X_T^{t,s}=e^{(r+\lambda-\frac{\sigma^2}{2})(T-t)+\sigma(W_T-W_t)}\left(s-\lambda\overline{S}_0e^{\mu t}\int_t^T e^{-(r+\lambda-\mu-\frac{\sigma^2}{2})(u-t)-\sigma(W_u-W_t)}du\right). \tag{28}$$

Remark 2 *Such a function u belongs to $C^{1,2}$ because we have a linear second-order PDE of the parabolic type, and arguing from the PDE theory, our payoff function is of the right type to ensure existence and uniqueness of the PDE solution which is $C^{1,2}$. Countable number of discontinuities in the payoff function do not matter.*

Proof: The PDE (26) is associated to the operator

$$\mathcal{L}f = \frac{\partial f}{\partial t} + \frac{1}{2}\sigma^2 s^2 \frac{\partial^2 f}{\partial s^2} + s(r+\lambda)\frac{\partial f}{\partial s} - \lambda \overline{S}_0 e^{\mu t}\frac{\partial f}{\partial s}$$

which generates the stochastic differential equation

$$dX_t = X_t((r+\lambda)dt + \sigma dW_t) - \lambda \overline{S}_0 e^{\mu t} dt, \; X_T = s,$$

the solution of which being:

$$X_T^{t,s} = e^{(r+\lambda-\frac{\sigma^2}{2})(T-t)+\sigma(W_T-W_t)} \left(s - \lambda \overline{S}_0 e^{\mu t} \int_t^T e^{-(r+\lambda-\mu-\frac{\sigma^2}{2})(u-t)-\sigma(W_u-W_t)} du \right).$$

For that, it helps to re-write $u(t,s)$ as an expectation:

$$u(t,s) = E[(X_T^{t,s} - K)^+],$$

where we obtain $X_T^{t,s}$ through a straightforward application of the Feynman-Kac's formula:

$$X_T^{t,s} = e^{(r+\lambda-\frac{\sigma^2}{2})(T-t)+\sigma(W_T-W_t)} \left(s - \lambda \overline{S}_0 e^{\mu t} \int_t^T e^{-(r+\lambda-\mu-\frac{\sigma^2}{2})(u-t)-\sigma(W_u-W_t)} du \right).$$

•

In other words, actually it is S_T before the first N jump time T_1 (9) but evolving from state s at time t. The removal of the "integro" component in PDE (14) is equivalent to the removal of jumps from the underlying SDE (8).

The integral in the above equation is proportional to the average of a geometric Brownian motion over time, i.e. when we manage the risk of price corrections, the risk premium depends on the full path of the Brownian motion on $[t,T]$, not just on its value at maturity. From the Finance point of view, it means that, when estimating the risk premium, one would not just care about where the stock price is going to end up at maturity (which is the case in the Black-Scholes world) but would also think about the average position of its price with respect to the fundamental value during the life time of the option. This reasoning suggests that, since the path-dependence is an important part of the risk management behaviour, it could be difficult to improve analytical tractability of the solution because it is the integral over time that introduces $\lambda \overline{S}_0 e^{\mu t} \partial_s u$ term into our PDE.

The most prominent case of path-dependent options are Asian options. Here, averaging over time is used to reduce the impact of the last minute price jumps on the final payout at maturity which actually is an example of a risk management behaviour but on the option buyer's side. In contrast, we are introducing a risk premium for the option sellers. Nevertheless, the mathematical formulation of the problems appears to be very similar: in both cases we have path-dependence and averaging over time.

4.1 Properties of the solution

The problem of solving PDE (26) is mathematically equivalent to the problem of pricing Asian options, and there are no known analytical time-domain solutions to it. Only the Laplace transform with respect to time is available in terms of hypergeometric functions [15]. Nevertheless, there are still some insights we can get from the equation itself. For example,

(25) yields $\partial_s C(t,s) = e^{-(r+\lambda)(T-t)}\partial_s u(t,s)$ and $\partial^2_{ss}C(t,s) = e^{-(r+\lambda)(T-t)}\partial^2_{ss}u(t,s)$. Then, we can differentiate the "expectation" form of $u(t,s)$ (27) with respect to s and obtain

$$\begin{aligned} \partial_s C(t,s) &= E[e^{-\frac{\sigma^2}{2}(T-t)+\sigma(W_T-W_t)}\theta(X_T^{s,t}-K)] \\ \partial^2_{ss}C(t,s) &= E[e^{(r+\lambda-\sigma^2)(T-t)+2\sigma(W_T-W_t)}\delta(X_T^{s,t}-K)], \end{aligned} \tag{29}$$

where differentiating the call option payoff gives us Heaviside θ and Delta δ functions. For the first result, we can commute expectation with differentiation because we can bound the non-negative expression inside expectation's brackets by the coefficient $e^{-\frac{\sigma^2}{2}(T-t)+\sigma(W_T-W_t)}$ which does not depend on s and has a finite expectation which is equal to 1. The result itself follows from the mean value and dominated convergence theorems.
Furthermore, the coefficient $e^{-\frac{\sigma^2}{2}(T-t)+\sigma(W_T-W_t)}$ can operate as a probability measure change (Girsanov's theorem) meaning that $\partial_s C(t,s)$ is in fact a cumulative distribution function of the following random variable:

$$Y_T^{t,K} = Ke^{-(r+\lambda-\frac{\sigma^2}{2})(T-t)-\sigma(W_T-W_t)} + \lambda\overline{S}_0 e^{\mu t}\int_t^T e^{-(r+\lambda-\mu+\frac{\sigma^2}{2})(u-t)-\sigma(W_u-W_t)}du, \tag{30}$$

and $\partial^2_{ss}C(t,s)$ is its probability density function:

$$\begin{aligned} \partial_s C(t,s) = E[\theta(s-Y_T^{t,K})] = P(Y_T^{t,K}\le s) \\ \partial^2_{ss}C(t,s) = \partial_s P(Y_T^{t,K}\le s) = f_{Y_T^{t,K}}(s) \end{aligned} \tag{31}$$

where actually $P_{|\mathcal{F}_t} = e^{-\frac{\sigma^2}{2}(T-t)+\sigma(W_T-W_t)}Q_{|\mathcal{F}_t}$.

Both derivatives are strictly positive suggesting that the call option's price is a strictly increasing and convex function with respect to the stock price. We can also see that $\partial_s C(t,s) = P(Y_T^{t,K}\le s)$ goes to 1 as s goes to infinity because it is a distribution function. The first derivative with respect to s denotes the position of the delta hedge, and it therefore makes sense for the value of the hedge to be an increasing function with respect to the stock price and not to exceed the stock price itself. These are essentially the same properties of the solution that we have in the Black-Scholes world. The convexity also means that, after a jump, the value of the replicating portfolio will always be below the call option's price because its value evolves along the tangent line which would be below the call option's price if it is a convex function of s. This means that jumps are always bad for option sellers, they present themselves as a risk, and one should therefore manage them as such.

It is also interesting to see what happens if we vary the payoff function which is equivalent to changing the terminal condition in the PDE for u (26). For example, we can make it equal to the stock price itself, and a search for a solution in the form of $se^{(r+\lambda)(T-t)}+f(t), t\to f(t), f(0)=0$ will give us:

$$u(t,s) = se^{(r+\lambda)(T-t)} - \frac{\lambda\overline{S}_0 e^{\mu t}[e^{\mu(T-t)}-e^{(r+\lambda)(T-t)}]}{\mu-r-\lambda}. \tag{32}$$

Then, we can plug the above expression for $u(t,s)$ into (25) and obtain the following trivial identity:

$$S_t = s. \tag{33}$$

The result is not as straightforward as it may seem, one needs to take several integrals to derive it but the equality nevertheless does hold which means that we can use the same pricing rule to price the stock itself.
It is also clear that we can apply our pricing rule to another trivial asset - a bond that pays one unit of "currency" at maturity. Setting the terminal condition in the PDE (26) to 1 gives us $u(t,s)=1$ and from (25)

$$P_{tT} = e^{-(r+\lambda)(T-t)} + \lambda e^{-r(T-t)} \int_t^T e^{-\lambda(T-v)} dv = e^{-r(T-t)}. \tag{34}$$

These two identities were certainly expected. It is obvious that we can use the risk-neutral measure Q to price the stock and a bond not just call options. They also imply that the call-put parity relationship between call and put option's prices (C_t, P_t) will be true as well:

$$C_t - P_t = S_t - Ke^{-r(T-t)}. \tag{35}$$

This is because the difference of the call and put payouts is a linear function of the stock price itself, and we can easily derive the above equation by taking expectations in our original pricing rule (12).

4.2 Boundary conditions

In order to solve PDE (26) numerically, it would be useful to have boundary conditions for large values of s and for $s=0$.

Proposition 4.3 *The boundary conditions for the functions u and C are:*
$u(t,s) = \partial_s C(t,s) = \partial_s u(t,s) = 0$ *for* $s \le 0$
and as a special case: $u(t,0) = \partial_s C(t,0) = \partial_s u(t,0) = 0$.

$$\lim_{s\to\infty} \partial_s u(t,s) = 0 \ ; \ \lim_{s\to\infty} \partial_s C(t,s) = 0.$$

Thus

$$u(t,+\infty) \approx se^{(r+\lambda)(T-t)} - K - \frac{\lambda \overline{S}_0 e^{\mu t}[e^{\mu(T-t)} - e^{(r+\lambda)(T-t)}]}{\mu - r - \lambda}$$

$$u(t,0) = \partial_s u(t,0) = 0, \tag{36}$$

and the expression for $C(t,s)$ (25) gives us:

$$C(t,+\infty) \approx s - Ke^{-(r+\lambda)(T-t)} - \frac{\lambda \overline{S}_0 e^{\mu t}[e^{(\mu-r-\lambda)(T-t)} - 1]}{\mu - r - \lambda} + \lambda e^{-r(T-t)} \int_t^T e^{-\lambda(T-v)} u(v, \overline{S}_0 e^{\mu v}) dv$$

$$C(t,0) = \lambda e^{-r(T-t)} \int_t^T e^{-\lambda(T-v)} u(v, \overline{S}_0 e^{\mu v}) dv \ ; \partial_s C(t,0) = 0. \tag{37}$$

Proof: Recall that $\partial_s u(t,s) = e^{(r+\lambda)(T-t)} \partial_s C(t,s)$ which is proportional to the hedge position ϕ_t. We do expect it to be zero because at $s=0$ the only risk we have is the risk of jumps, and we are not hedging for it.

Indeed, let us recall the solution exhibited in (28):

$$X_T^{t,s} = e^{(r+\lambda-\frac{\sigma^2}{2})(T-t)+\sigma(W_T-W_t)} \left(s - \lambda \overline{S}_0 e^{\mu t} \int_t^T e^{-(r+\lambda-\mu-\frac{\sigma^2}{2})(u-t)-\sigma(W_u-W_t)} du \right),$$

and note that, since the process $\lambda\overline{S}_0 e^{\mu t}\int_t^T e^{-(r+\lambda-\mu-\frac{\sigma^2}{2})(u-t)-\sigma(W_u-W_t)}du$ is continuous and strictly increasing, $X_T^{t,s}$ shall take negative values after time

$$(38)\qquad \tau := \inf\{T \geq t : \ \lambda\overline{S}_0 e^{\mu t}\int_t^T e^{-(r+\lambda-\mu-\frac{\sigma^2}{2})(u-t)-\sigma(W_u-W_t)}du = s\} < \infty.$$

Moreover, the process will take strictly negative values at maturity if it starts from zero or below zero. That however means that expectations in the expressions for $u(t,s)$ (27) and $\partial_s C(t,s)$ (29) are zeros because the payoff and Heaviside functions are zeros for negative values of their arguments, and we can immediately see that $u(t,s) = \partial_s C(t,s) = \partial_s u(t,s) = 0$ for $s \leq 0$ and as a special case: $u(t,0) = \partial_s C(t,0) = \partial_s u(t,0) = 0$.

As for the boundary condition for large values of s, we can derive it from PDE (26) by replacing its terminal condition with $u(T,s) = s - K$. In fact, we can use the solution that we had obtained for the terminal condition $u(T,s) = s$: (32). Since PDE (26) does not have $u(t,s)$ but only its derivatives, we only need to reduce it by K:

$$(39)\qquad u(t,+\infty) \approx s e^{(r+\lambda)(T-t)} - K - \frac{\lambda\overline{S}_0 e^{\mu t}[e^{\mu(T-t)} - e^{(r+\lambda)(T-t)}]}{\mu - r - \lambda}.$$

We also want to check that the following asymptotic expression holds:

$$(40)\qquad |s e^{(r+\lambda)(T-t)} - K - \frac{\lambda\overline{S}_0 e^{\mu t}[e^{\mu(T-t)} - e^{(r+\lambda)(T-t)}]}{\mu - r - \lambda} - u(t,s)| \to 0$$

as $s \to +\infty$. It helps to re-write it as

$$(41)\qquad |E[(X_T^{t,s} - K)] - E[(X_T^{t,s} - K)^+]| = E[(K - X_T^{t,s})^+]$$

which tends to zero as $s \to +\infty$ because it is positive and a strictly decreasing function of s for $s > 0$. The later is because the above expectation solves for the put option, and the put-call parity relationship (35) suggests that the put option's derivative with respect to s is in $(0,1]$ and tends to zero when $s \to +\infty$ which is what what we would expect in the Black-Scholes "world". This concludes the proof. •

It is interesting that, according to Remark 1, we can restrict PDE (26) to $s \in R^+ - \{0\}$ and forget about negative values of the stock price. Note that this is a very special case which comes from the fact that the process X never goes back to positive values after it touches the line $s = 0$.

4.3 Solving the PDE numerically

Since our problem (26) is mathematically equivalent to the problem of pricing Asian options, one may refer to an abundance of numerical methods available in this domain. Unfortunately, there are no analytical solutions which significantly undermines practical applications of our model. However, it is helpful to note that, since we are trying to explain the implied volatility phenomena, our objective is to adjust the Black-Scholes formula rather than to come up with fundamentally different solutions. It is therefore natural to consider asymptotic methods, and, in particular, those that have something like the Black-Scholes formula as their first approximation. The β expansion asymptotic method [14] developed for the purpose of pricing a floating-strike nicely fits this criteria. It is an example of the matched asymptotic expansions method which has been first introduced into the field of financial mathematics by [9], and it is an extension of the work by [5] which derives Black-Scholes-like asymptotic solutions for the "Asian" PDE.

From now on, t is fixed, so is $S_t = s$, and $\overline{s} = \overline{S}_t$. Along with μ, r, σ, λ, K and T, these now become inputs into our pricing algorithms. The following result is proved in [14].

Proposition 4.4 *The time t being fixed and $\tau = T - t$, the solution of a floating-strike put Asian PDE can be reduced to the solution of the following PDE:*

$$\partial_\tau \psi(\xi,\tau,\rho,T,\lambda) = \frac{1}{2}(\xi - \frac{1}{\rho T}(1-e^{-\rho\tau}))^2 \partial_{\xi\xi}\psi(\xi,\tau,\rho,T,\lambda) \; ; \; \psi(\xi,0,\rho,T,\lambda) = (\xi-\lambda)^+ \tag{42}$$

with some real constants ξ, $\tau \geq 0$, ρ, $T > 0$, and $\lambda > 0$. Thus we can reduce the above PDE (42) for pricing PDE (26) through a straightforward change of variables. The solution satisfies:

$$u(t,s) = \psi\left[\zeta\left(s,\overline{s},\mu,\lambda,\tau,\sigma,K,r\right), \sigma^2\tau, \frac{\mu-r-\lambda}{\sigma^2}, -\frac{\sigma^2}{\lambda\overline{s}e^{\mu\tau}}, K\right], \tag{43}$$

where the function ψ is the same as in [14], ζ is defined as a natural Log function similar to the one in the Black-Scholes formula:

$$\zeta\left(s,\overline{s},\mu,\lambda,\tau,\sigma,K,r\right) = Log\left[\frac{s}{K}e^{(r+\lambda)\tau} - \frac{\lambda\overline{s}e^{\mu\tau}\left(1-e^{-(\mu-r-\lambda)\tau}\right)}{K(\mu-r-\lambda)}\right] + \frac{1}{2}\sigma^2\tau, \tag{44}$$

and the expression for ψ also looks like the Black-Scholes model:

$$\psi(\zeta,\tau,\rho,L,K) = e^{\zeta-\frac{\tau}{2}}K\Phi_G\left[\frac{\zeta}{\sqrt{\tau}}\right] - K\Phi_G\left[\frac{\zeta-\tau}{\sqrt{\tau}}\right] + \epsilon(\zeta,\tau,\rho,L,K), \tag{45}$$

where Φ_G denotes the cumulative normal distribution function, and ϵ is an asymptotic expansion of the remainder:

$$\epsilon(\zeta,\tau,\rho,L,K) = \frac{1}{\sqrt{2\pi\tau}}\left(\frac{\tau}{L}\right)e^{-\frac{\zeta^2}{2\tau}}\sum_{n=1}^{n_\beta}\sum_{i=0}^{n_\xi}\sum_{j=0}^{n_\tau}\sum_{k=0}^{j-1}\frac{b_{n,i,j,k}\zeta^i\tau^j\rho^k}{n!(KL)^{n-1}}. \tag{46}$$

The real-valued constants $\{b_{n,i,j,k}\}$ can be obtained from the algorithm described in [14].

In the spirit of the matched asymptotic expansions method, the solution is obtained as a closed form expression suitable for all practical purposes. We should note that there is no proof for the approximation to converge, and we should not deviate too much from the Black-Scholes formula. The method does however make the model usable for all practical purposes because standard methods for solving the problem of Asian options may not be fast enough.

As we can see, the algorithm for $u(t,s)$ is a relatively simple analytical function including a cumulative normal distribution function and finite sums but, in order to obtain the function C (25), we need to take a finite one-dimensional integral over time. The integrand however is a continuous, strictly increasing function with no singularities on $[t,T]$, and we can therefore easily approximate it with a sum.

While the asymptotic method gives us an essentially analytical solution to our problem, criteria for its convergence are not known and it is therefore necessary to benchmark its performance against a standard algorithm used to solve this particular problem. For this purpose, we used the finite difference method with an implicit finite difference scheme [3].

In summary, we found that, if we are to consider high ($\lambda \geq 3$) Poisson intensities for price corrections, the asymptotic algorithm does not perform well. Large deviations from the Black-Scholes world are not desirable. Another critical parameter is the time to maturity $\tau = T - t$ which works together with the volatility σ (the critical input is $\sigma^2\tau$), the algorithm gives better accuracy for lower volatilities or/and when closer to maturity.

It is interesting to see what happens when we get close to the borders of the allowable ranges and see how errors and numerical instabilities arise. Figure 1 demonstrates what goes wrong with the asymptotic expansion algorithm. When going far enough from maturity we see a steady increase in the absolute error slowly oscillating along strikes. The error will start developing closer to maturity if we increase λ or/and σ but the shape will be similar to what we see in the figure. As for the finite difference scheme itself, it does contribute to some errors but they will reduce if we are to reduce the size of steps in the mesh. In particular, the “ridge" along the strike of around 100, for both $u(t, s)$ and option’s Delta, is because the mesh steps are not small enough.

The asymptotic expansion algorithm is therefore good for all practical purposes. This is for calculating both option’s price and Delta, one just needs to treat results as estimates if we have unusually high Poisson intensities (more than 3) or/and long maturities (more than 6 months).

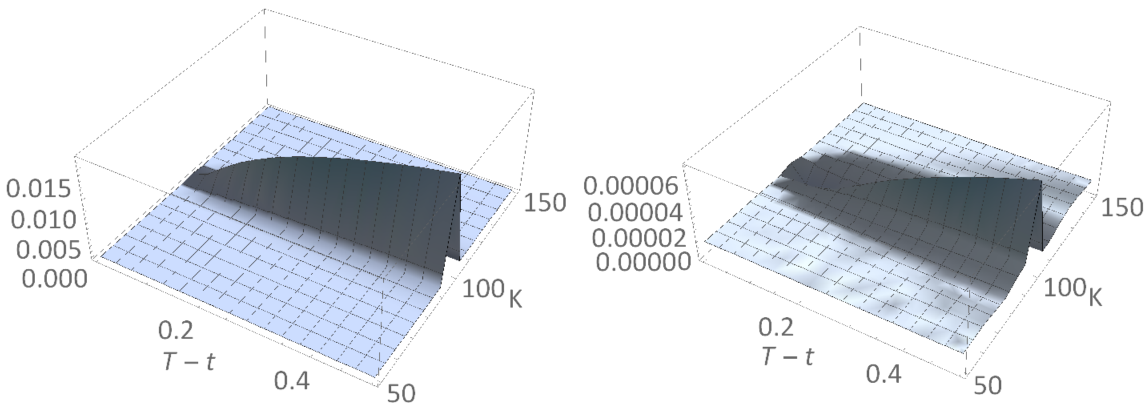

Figure 1: Finite difference vs. Beta expansion algorithm: difference for $u(\tau, s)$ on left hand and option’s Delta on right hand as functions of $\tau = T - t$ and K; $\sigma = 0.2$, $\mu = 4.125\%$, $r = 0.15\%$, $\lambda = 0.25$, $\overline{s} = 100$, $s = 100$.

4.4 Calibration algorithms

Having algorithms for the option’s price may not be sufficient in itself. It is often that we want to solve a reverse problem - what are the parameters implied by the model given historic prices? This is particularly important for our problem because one of our objectives is to have a mechanism to measure the probability of price corrections λ implied by the market. The process for solving this problem is usually referred to as calibration. We need to calibrate some of our input parameters (e.g. σ, $\overline{s}$, λ and may be μ) so that to achieve the best possible fit (in the least squares sense) of the model's output to historical data, and there are many guides for practitioners on how to do it, e.g. [4]. We can either calibrate to prices or implied volatilities. The latter approach is preferable because the volatility smile and skews are the manifestation of the markets aversion to risk, and we can guess some parameters from the shape of the volatility surface itself, e.g. it will tell us where we are with respect to the fundamental value of the stock - above, below or approximately at it. There is only one thing that we will improve on to take in account the liquidity - we will introduce weights to implied volatility errors by dividing them with the difference in option's bid and ask prices. This approach is often used in the industry because the weights amplify errors for more liquid options thus skewing the result towards options that are more actively traded.

As for the optimization procedure itself, one of the popular algorithms arising in this context is the Levenberg-Marquardt algorithm [12] which is an example of an iterative algorithm for finding a local minimum in a multi-dimensional setting.

5 Numerical results

The objective of this section is to test our model with the real markets data. If the volatility surface is formed in response to the risk of price corrections, we should expect the model to produce similar shapes of the volatility surface (like skews and smiles) as observed in reality. Moreover, it should fit actual data points with reasonable values for the three new (vs. Black-Scholes) parameters: $\overline{s}$, λ and μ. The latter is to make sure that we can interpret them as market's perceptions about the fundamental value of the stock $\overline{s}$, intensity of price corrections λ, and expected rate of return on the stock μ.

We start by showing how well we can fit the model to data by comparing its performance with the Heston model and then continue to explore ranges of parameters obtained through calibration.

5.1 Calibration to market

Stocks from the NASDAQ 100 index are highly liquid, and standard American call and put contracts are traded at many maturities and strikes. We look at one month worth of option data for randomly selected ten stocks from the index. Our primary objective is to calibrate the model's outputs to implied volatilities. After all, we assume that the shape of the implied volatility surface comes from the markets aversion to risk, and it is this risk that we want to price. In all instances, we used the finite difference method because it is guaranteed to converge with no numeric instabilities, and we have to allow for a more general case of American options with discrete dividends.

When calibrating to data, we learned several things that defined our calibration strategy.
• Firstly, when calibrating to contracts with different dates, we could get relatively stable values for σ but not for $\overline{s}$. They reflect market's perceptions about the fundamental value of the stock and Poisson intensity of price corrections towards this value, and these perceptions are just not stable over time. The "bullish" market in the underlying may become "bearish" in an instant after a new information is released to the market. We do expect these parameters to be piecewise constant (over quiet, no new information periods in the market activity) which will still make the analytics hold. However, since it is hard to know if the time period was really "quiet" and, just for the sake of simplicity, we decided to calibrate to only one volatility shape, i.e. to contracts with the same stock symbol, date, time to maturity and the option type (call/put). One month worth of data over ten randomly selected stocks gave us 472

calibration results across 6,238 contracts.
• Secondly, we found that the calibration could be very slow and may result in unreasonable values for μ, especially when the stock is not priced on value ($\bar{s}$ is far from s). We hypothesized that it must be because the model's outputs have low sensitivity to μ, and it makes sense to set it to a constant (let us say 4.125% per year). As well as making the calibration process faster, it has allowed to remove one degree of freedom and use the following expression for the Standard Estimation Error:

$$SEE = \sqrt{\sum_{i=1}^{N} \frac{([\text{impl. vol. model}]_i - [\text{impl. vol. data}]_i)^2}{N-3}}, \tag{47}$$

where [impl. vol. data] is obtained from the market's data (it is volatility implied by the Black-Scholes formula and market prices), [impl. vol. model] is a result of the calibration algorithm (it is volatility implied by our model and market prices), and the sum is over N data points. This is a standard expression for the SEE with three degrees of freedom (σ, $\bar{s}$, λ), and it is the first thing that differentiates our model from the Heston model. The latter has five degrees of freedom, and, when measuring its goodness of fit, 5 has to be used instead of 3.

• We also found that calibration results could be sensitive to the initial value of $\bar{s}$, the value from which the calibration process starts. There are presumingly multiple minima because large corrections with low probabilities demand similar price as small corrections with high probabilities. We have therefore calibrated with three different starting points for $\bar{s}$ (50%, 100% and 150% of the stock price) and selected best optimization results with respect to SEEs.

5.2 Performance vs. the Heston model

We found that, from the goodness of fit point of view, the comparison with the Heston model is not very interesting. As the box and scatter plots from the Figure 2 show, the models perform very similarly, and the errors are strongly correlated. We can obtain a very good linear regression model fit ($R^2 = 0.50$), and the correlation coefficient is close to one (0.78).

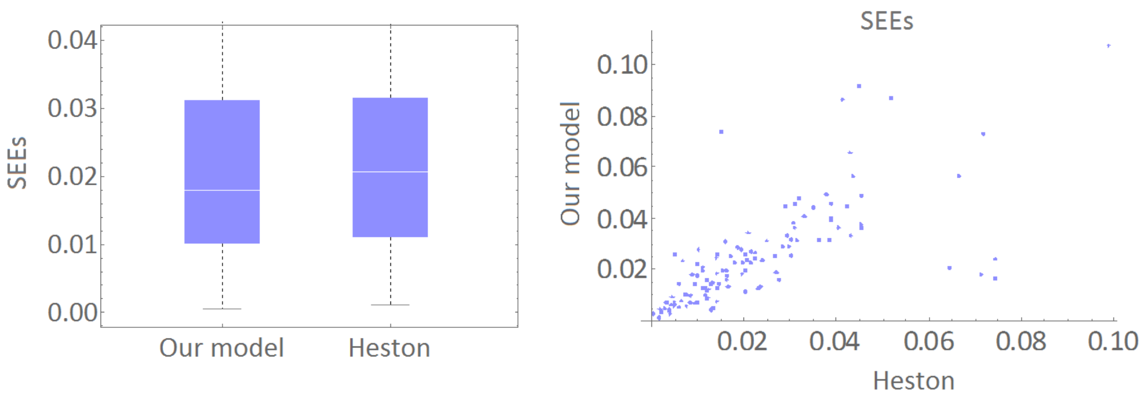

Figure 2: Standard Estimation Errors (SEEs) for my model (SS) vs. Heston model, box and scatter plots for SEEs.

It is still true that models may perform differently for a certain range of parameters, and we indeed see that the model does seem to perform slightly better when the contracts maturity is less than three weeks. However, the p-value for the hypothesis that there is no difference in the means of SEEs (t-test) is 0.18 which is too high to claim the result as statistically significant. We also find no statistically significant and/or interesting differences in the means of SEEs when looking across other parameters including stock symbol and the type of the volatility surface involved (smile vs. skews, where we define smile as $0.9 \leq \bar{s}/s \leq 1.1$).

5.3 Volatility σ

The meaning of the parameter σ in our model is no different to that in the Black-Scholes model. It is a multiplier in front of the Brownian motion, a scaling parameter that defines the amplitude of noise driving the stock price. We obtain it through calibration but it can also be estimated through the statistical properties of log returns - we expect them to be normally distributed with the standard deviation defined by the time scaled σ. That still holds true in our model except that we need to exclude jumps from the time series. The jumps are not part of the noise, they are corrections in response to a new information released to the market. We therefore expect the range of σ to be similar to the range of implied volatilities, and this is indeed is the case. Looking at Figure 3 we can see that for all stocks the median is slightly below the median of implied volatilities, and for all stocks we see a reduction in the data range for σ.

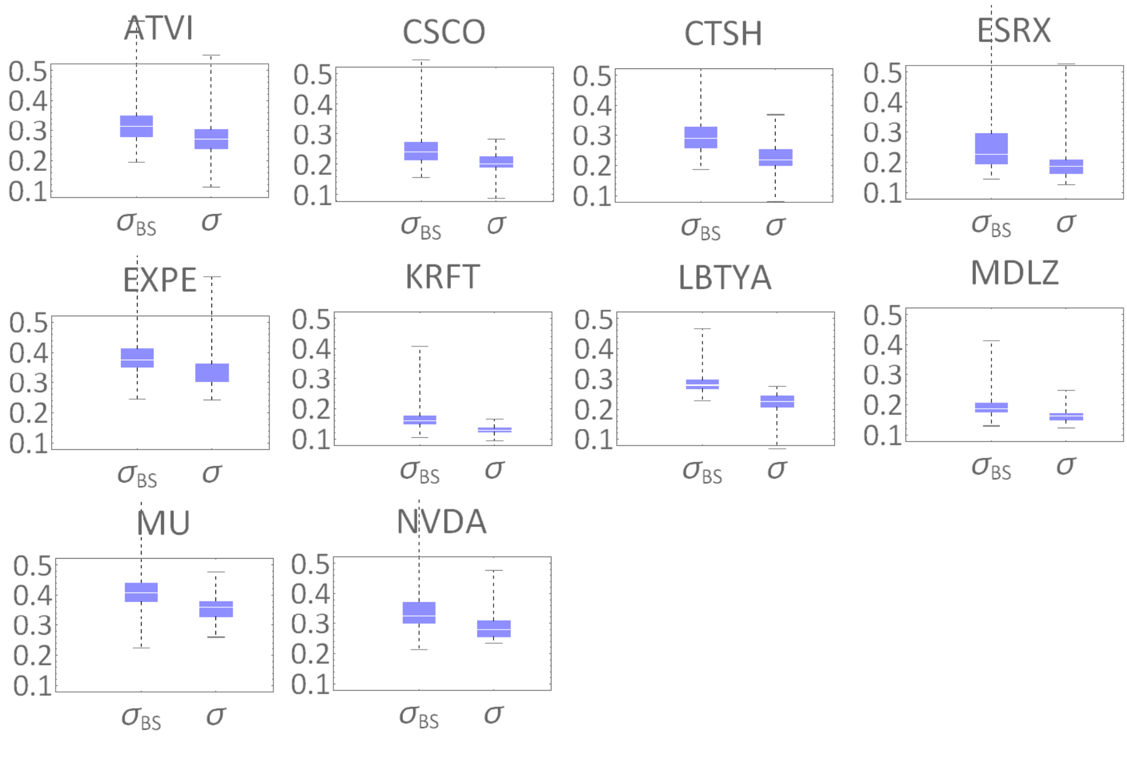

Figure 3: Box plots for implied volatilities (σ_{BS}) left hand vs. volatilities from our model (σ) right hand.

The result was almost certainly expected. In the Black-Scholes model, σ_{BS} is slightly inflated because it needs to account for the risk premium of price corrections and, by fitting the volatility surface with our model, we significantly reduce daily variations in the parameter resulting in the data range being smaller over time.

We can also see that, for the purpose of calibration, it is very easy and intuitive to guess σ. When we look the shape of an implied volatility curve (e.g. as a function of strikes), we know that σ should be slightly less than all implied volatilities. This is very much unlike the Heston model where σ is a volatility of the square of the volatility. It has a very different meaning, and one can not guess it as easily as σ in our model.

5.4 Fundamental value of the stock $\overline{s}$ and λ

The remaining two parameters are the most interesting. They are new (vs. Black-Scholes model), and they are all about the risk of price corrections. Let us start by looking at the ranges of optimum values that we have managed to achieve through calibration. Figure 4 shows histograms for calibrated $\overline{s}/s$ and λ.

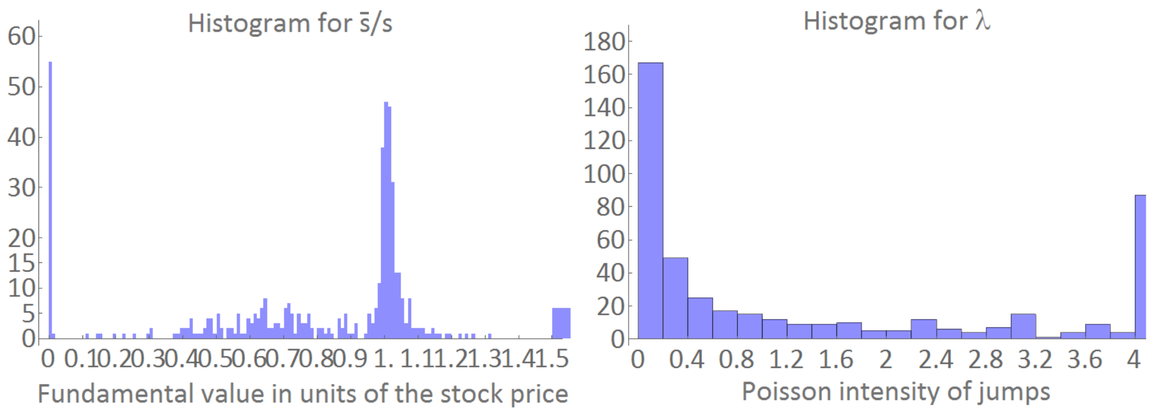

Figure 4: Histograms for $\overline{s}/s$ and λ.

First thing we notice is that we have three distinct clusters for the fundamental value of the stock. We have a spike below 10% (close to zero, not immediately visible), in the range of $90 - 100\%$ and something could be going on above 150% of the stock price. The first cluster is when markets are expecting the stock to default, second is the implied volatility smile region where the stock is perceived to be priced on value, and the right cluster is when an extreme upward correction is expected. For the remaining values, we have a skew that is not well pronounced and this is a grey region where markets are undecided if there should be an extreme movement or if the stock is priced on value after all. These are market behaviours implied by the model.

We also have the Poisson intensity of price corrections λ which is concentrated around smaller values, in fact, its median is around 0.5 meaning that in approximately half of the cases markets expect less than one correction in two years on average. This is a very desirable result because high values for λ will take us outside of the standard risk-neutral valuation framework and may also result in numeric instabilities for the asymptotic algorithm.

The low values of λ do undermine the motivation behind the Heston model which relies on local changes in volatility to explain the implied volatility surface. Fundamentally, jumps with the Poisson intensity of less than 0.5 can not be localized to any short term statistic of

Table 1: Calibration results for the ranges of $\overline{s}/s$ and λ

	low intensity $\lambda \leq 0.5$	high intensity $\lambda > 0.5$
extreme upward correction is expected, $\overline{s}/s \geq 1.5$	6	0
the stock is priced on value, $0.9 \leq \overline{s}/s \leq 1.1$	29	212
default is expected, $\overline{s}/s \leq 0.1$	56	0
undecided $0.1 \leq \overline{s}/s \leq 0.9$	141	28

the stock price. They are about the fear of price corrections, not about the price corrections themselves.

It is also interesting to see how the ranges of $\overline{s}/s$ and λ are related. As Table 1 shows, when the stock is priced on value (i.e. when we have the implied volatility smile, second line), the more often cases are in case of a high intensity for Poisson jumps ($\lambda > 0.5$). While as the default and extreme upward correction cases are concerned (lines 1 and 3) - all results have calibrated to some λ less than 0.5.

There are not many results related to the extreme upward correction case (only six), markets seem to have rarely expected this to happen. However, it is reasonable to assume this region to be a symmetrical reflection of the default case and lead to similar findings as the default case.

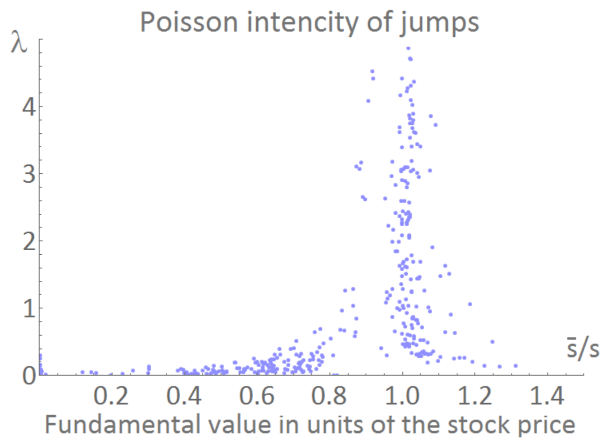

Figure 5: λ versus $\overline{s}/s$.

We can see the dependency between $\overline{s}/s$ more clearly on figure 5. When stock is perceived as priced on value, we get a spike in the intensity of jumps. This is reasonable, when λ is high, jumps are more likely, and having large $|\overline{s}/s - 1|$ would lead to larger perceived profits for technical traders holding positions in stock. This will create higher demand or supply and reduce $|\overline{s}/s - 1|$. On the other hand, if λ is small, there is not too much jump risk, and markets can afford higher misalignment between the stock price and it's fundamental value.

The dependency may also undermine our assumption that λ and $\overline{s}$ are constants - they are not, they change over time. In order for the analytics to hold, we have to assume that these parameters are piecewise constant. This is a reasonable assumption because we expect λ and $\overline{s}$ to change only when there is a new information released to the market. These parameters are constant as long as we observe the same implied volatility shape but the change in the implied volatility shape means that market's perceptions about λ and $\overline{s}$ have changed and we have to recalibrate our model accordingly. That, for example, means that pricing derivatives in an environment when default is expected (implied volatility skew) is not the same as pricing when stock is perceived to be price on value (implied volatility smile).

5.5 Summary of numerical results

We find that there is no evidence to suggest that our model performs better or worst compare to the Heston model. At the same time, we only have three degrees of freedom (σ, $\overline{s}$, λ) vs. five in the Heston model. Moreover, we know that we can often estimate $\overline{s}$ and σ with a reasonable accuracy just by looking at the shape of the implied volatility surface. For example, if we have an implied volatility smile, it should be centered around $\overline{s}$, and we can expect σ to be somewhere just below the minimum implied volatility.

We have also found that, since in half of the cases we have calibrated to some λ less than 0.5, it may not be efficient to explain implied volatility phenomena through a local behaviour of the stock (e.g. by introducing a stochastic volatility). Fundamentally, the phenomenon does not arise from the stock dynamics but from the fear that price corrections may happen. Let us now move on to see how we can apply our model in the real world of pricing derivatives.

6 Practical applications

The objective of this section is to demonstrate how we can apply our model in the real world of pricing derivatives. One can assume that the applications would be similar to the models like the Heston model, and this is indeed is true. One of the main purposes of our research was to introduce a new model for pricing derivatives. We just need to be aware that we are pricing a risk. Firstly, this is not a precise science (probabilities of rare events are hard to estimate from data) and, secondly, pricing risks differently from competitors (assuming different intensity parameters λ) may not align with a business strategy. Overpriced risks means derivatives would not sell and undercutting competition by taking higher risks may not be the right thing to do either. It is therefore important that one may not wish to estimate λ and $\overline{s}$ from things like the history of price corrections and the fundamental analysis. This application of our model would be too simplistic.

In reality, we would still want to calibrate to data, and note that it is sufficient to have just one volatility smile or skew for that purpose. Then we can estimate our risk parameters (λ, $\overline{s}$) and use this information to make more informed decisions about managing the risk of price corrections, i.e. we would not use the model to introduce different prices but rather to manage risks associated with underpricing or overpricing the risk of price corrections by the market. For example, if markets are pricing derivatives within the default region ($\overline{s}/s \leq 0.1$), which incidentally covers 12% of our calibration results, we can estimate the Poisson intensity of default implied by the market (for our results it is 0.05 on average) and compare it to the Poisson intensity of default implied by credit derivatives markets. Any difference in these parameters may lead to more efficient risk managing strategies or even a possibility of an arbitrage. On the other hand, if the market is assuming that the stock is priced on value (51% of our results), we could be underpricing risks by ignoring possible corrections in the stock price or/and default. In an ideal world, that would lead to more investments in risk management strategies that are mitigating this type of risk.

However, financial institutions may choose not to do so which leads us to the second application of our model - measuring the price of risk implied by the market and, if excessive risks are being taken, regulating these markets, e.g. by requiring derivatives sellers to disclose risk parameters or even hold reserves calculated from them. We are not only talking about call and put options but about much wider class of derivatives. The model can be extended to any type of the payoff function at maturity. We can therefore use it as a tool to identify "toxic" derivatives, derivatives where we do not have λ explicitly included in the price and where the financial impact of the mispricing of risk is at its highest. That would be helpful for formulating regulation strategies for the financial industry as a whole.

Finally, we would like to cover one of the most interesting but perhaps less scientific applications of our model. One may hope that having estimates for $\overline{s}$ and λ, as implied by the market, may help to predict where the stock price will go next, e.g., if we are in the default region, we might want to bet on the stock price going down. Hedge funds, for example, would be interested in this because making money on predicting future stock movements is not unusual in this industry.

Let us now look at the price of one of our stocks (ATVI) and mark days when the market is expecting it to default (red dots), when the stock is perceived as priced on value (green dots), and when the market expects an extreme upward correction in the price (yellow dots). The results are shown in Figure 6.

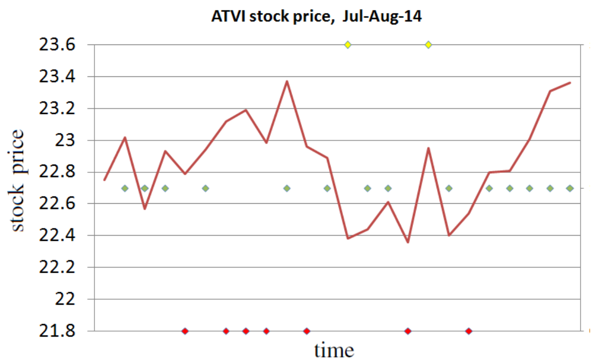

Figure 6: ATVI stock price, 15th July 2014 - 15th August 2014. Red dots denote days with a default expected, greens are priced on value days and yellows are where we expect an extreme upward correction.

We can see that for most days the stock is being perceived by the market as priced on value but there are seven days when markets were expecting the stock to default and two days when an extreme upward correction was expected. So, the question therefore is whether markets can predict future stock price movements. Interestingly, we do see that the largest downward and upward movements follow a cluster of red dots and two yellow dots respectively. Unfortunately, the result has no statistical significance but is more of a curiosity. One can only hope that market's perceptions about the true value of a stock has anything to do with the real value implied by the company fundamentals. This is however is a separate topic, outside the scope of this research.

7 Conclusion

We introduced a derivatives pricing model that interprets implied volatility smiles and skews as the market's aversion to the risk of price corrections. The model takes three new parameters (vs. the Black-Scholes): the fundamental value of a stock $\bar{s}$, the Poisson intensity of price corrections λ, and expected rate of return on a stock μ. Similarly to the Black-Scholes parameters r and σ, the three new ones can be interpreted and estimated independently of the model.

We introduced the fundamental value of a stock as a deterministic, exponential function, and price corrections as Poisson jumps to this function. Surprisingly, that took us on a journey beyond the standard risk-neutral valuation framework, and we had to go to the very basics of financial mathematics in order to understand what was going wrong. We demonstrated that, while it is possible to have an alternative formulation for the model from within the standard risk-neutral valuation framework, fundamentally, we are not operating in a risk-neutral environment, the model after all is about pricing risk.

Another obstacle was the technical complexity of the pricing algorithms that arose in this context. It turned out that pricing the risk of price corrections is similar to pricing Asian options - we need to take expectations over a geometric Brownian motion integrated over time. Since having no analytical formulae for the solution significantly undermines most practical applications, we introduced an essentially exact asymptotic algorithm which is simply an adjusted Black-Scholes formula plus a remainder calculated with finite sums.

We have test-driven our model for ten stocks randomly selected from NASDAQ 100 index. The model seems to perform no differently to the Heston model but, what is different is that one can estimate its parameters from the parameters implied by the Black-Scholes model and from the shape of an implied volatility surface. The volatility σ is just slightly lower than implied volatilities, and, similarly to the Black-Scholes formula, μ can be ignored, we can just set it to a constant. As for the parameter $\bar{s}$, it defines the minimum of the implied volatility surface, and we should expect λ to be less than 0.5 if we see a skews and to be above 0.5 in case of a smile.

As for the practical applications of our model, we believe they go beyond the traditional derivatives pricing role. We can measure λ and $\bar{s}$ implied by the markets across many derivative types and assess if markets are over or underpricing risks related to price corrections in the underlying. We can then identify "toxic" derivatives, derivatives where we do not have λ explicitly included in the price, and where the financial impact of the mispricing is at its highest. That would clearly be helpful in formulating regulation strategies for the financial industry as a whole. As the events of the 2008 financial crisis demonstrated, it is hard to regulate markets with λ explicitly included in the pricing algorithm (credit derivatives), let alone derivatives where λ does not even appear in the pricing formula.

While the mathematical community talks about risk-neutral pricing, in practice, there are some real sources of risk, and the possibility of price corrections is one of them. At the same time, given the size of the derivatives industry, even small mispricing of risk may lead to significant losses for the system because everybody calibrates derivatives pricing models to historical prices which means pricing risks at the same level as everybody else does it. We genuinely believe that our research has a very high value to the society as an enabler for stabilizing financial markets and preventing the next financial crisis.

The next steps should therefore most certainly include testing applications of our model beyond standard vanilla options and quantifying systematic risks in the derivatives industry. Moreover, since we can introduce mean-reverting jumps to any process, we can go beyond equity markets and expand applications to virtually any type of derivatives product.

Acknowledgements

We would like to thank Prof. William Shaw for his guidance and support. The constrictive feedback he gave during our meetings and presentations was invaluable in taking this work to its logical conclusion. We would definitely like to acknowledge contribution from the examiners during MPhil-PhD transfer (Prof. Lane

Hughston) and PhD viva (Prof. Sam Howison, Dr. Andrea Macrina). We have also to thank Pr Dominique Lépingle for his valuable help, and the ones who listen to us in Séminaire Bachelier, helping us with their remarks, comments and questions.

References

[1] Bellamy N. and Jeanblanc, M. 2000. *Incompleteness of markets driven by a mixed diffusion.* Finance and Stochastics, 4(2), 209-222.

[2] Black F. and Scholes M. 1973. *The pricing of options and corporate liabilities.* The journal of political economy, 81(3), 637-654.

[3] Crank J. and Nicolson, P. 1947. *A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type.* In Mathematical Proceedings of the Cambridge Philosophical Society, 43(1), 50-67. Cambridge University Press.

[4] Crosby J., and Nicolson P. 2010. *Practicalities of pricing exotic derivatives.*.

[5] Dewynne J. N. and Shaw W. T. 2008. *Differential equations and asymptotic solutions for arithmetic Asian option: "Black-Scholes formulae" for Asian rate calls.* Euro Jnl of Applied Mathematics.

[6] Gatheral J. 2000. *The Volatility Surface.* Wiley Finance.

[7] Geman H. and Roncoroni A. 2006. *Understanding the fine structure of electricity prices.* The Journal of Business, 79(3), 1225-1261.

[8] 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(2), 327-343.

[9] Howison S.D. 2005. *Matched asymptotic expansions in financial engineering.* J. Engrg. Math. 53 (2005), no. 3-4, 385–406.

[10] Jeanblanc M., Yor M., Chesney M.: Mathematical Methods for Financial Markets. Springer, Berlin, Heidelberg, New York (2009).

[11] Lépingle D. et Mémin J. 1978. Sur l'intégrabilité uniforme des martingales exponentielles, Z. Wahrscheinlichkeitstheorie v. Gebiete 42, 175-203.

[12] Levenberg K. 1944. *A Method for the Solution of Certain Non-Linear Problems in Least Squares.* Quarterly of Applied Mathematics, 2, 164-168.

[13] Merton R. C. 1974. *On the pricing of corporate debt: The risk structure of interest rates.* The Journal of Finance, 29(2), 449-470.

[14] Siyanko S. 2012. *Essentially exact asymptotic solutions for Asian derivatives.* European Journal of Applied Mathematics, 23(03), 395-415.

[15] Yor M. 1992. *On some exponential functionals of Brownian motion.* Adv. Appl. Prob. 24: 509-531, 1992a.