LELIEVRE Clément

28612

Master 2 thesis

Research paper superviser : Mr Philippe HENROTTE

HEC PARIS

JUNE 2014

Pricing Down&In puts

Abstract

In this paper, we look to price down&in puts, especially in the case of an assymetric volatility smile and with a focus on symmetry method. The aim is to assess the extent of the discrepancy in pricing relative to the price obtained when using a symmetric volatility smile.

We first define this option, and price it using Monte-Carlo, PDE, image methods and lastly closed-form formulae in the Black&Scholes context ; we then build an asymmetric volatility smile by adding jumps and stochastic volatility and re-price the option. Lastly, we compare the prices obtained and try to establish the robustness of the symmetry method in this second context.

I Introdution

What barrier options are

A down&in (D&I) put is an exotic variant of a standard (plain vanilla, european) put : it is a put that activates, i.e. starts existing as a standard put, if and only if the underlying asset price crosses a predefined threshold (called the barrier) by above, at anytime during the life of the option.

So its payoff at maturity is defined by: $\max(0;K-S_T)*\mathbf{1}\{\min S_t <= B\}$ where :

- K is the strike
- S_t is the spot at time $t \in [0;T]$
- $\mathbf{1}\{A\}$ is the indicator function of event A : worth 1 if A happens, else 0
- B is the barrier

From this definition, one can immediately notice that obviously the barrier must be set below S_0 (the underlying spot price at inception), else the D&I put is nothing else but a standard put. We also see that studying the law of the minimum of the underlying asset is a natural way of studying this type of option.

We say that this put has an american barrier whereas the exercise-style is european (the latter can be american, but we will focus on european-style options here), because the barrier can be breached at any time between inception and maturity, but exercise can only happen at maturity. The D&I put belongs to the first-generation exotic options, which started being traded in the 1990s. Then all sorts of variants emerged, only bounded by imagination, for example we started to see traded on a regular basis double barrier options in a context of emergence of structured products tailored for clients. For instance, a double knock-out option shares similarities with a short straddle position in terms of market view.

An important question for investors in barrier options is the frequency at which the spot price is observed, as it impacts the likelihood of the option knocking in/out. Broadie, Glasserman and Kou provide a means of adjusting the closed-form formula for cases where price observations are discrete, in « A continuity correction for discrete barrier options » which we will use in our pricing.

Another issue is the potentiality of price manipulation, especially regarding underlying assets of poor liquidity, against which some second-generation options like the « Parisian » options were designed : their payoff has more constraints relative to a barrier option in that the underlying price must remain below/above the barrier for a certain time period (either cumulative or in a row) in order for the option to activate/remain active, while a barrier option does not impose any conditions regarding the time during which the barrier must be breached. Therefore, intuitively « in » parisian options are cheaper than « in » barrier options and « out » parisian options are more expensive than « out » barrier options.

The down&in put is, among the eight possible first-generation barrier options (call/put, down/up, in/out), more traded than most others, usually by very bearish investors and/or by investors looking for an insurance against a market crash.

In practice, « in » options rarely have cash rebates so we will assume the rebate is null in this paper.

Why they exist

As highlighted by E.Derman in « The ins and outs of barrier options », there are three main reasons to use barrier options instead of standard ones :

1) They may match investors' anticipations of market behaviour more closely than standard options
   Consequently, they enable an investor to avoid paying for scenarios which he thinks are very unlikely. In our case, a D&I put spares the cost of all the scenarios where the spot over the life of the option on average moderately decreases (without crossing the barrier), i.e. scenarios in which a standard put of same features would end up in the money (ITM) while a D&I put would expire worthless.

2) They may match hedging needs more closely
   Suppose you own a given stock and you set a stop loss whenever the stock loses 10% relative to the price you bought it at. A down&out (D&O) put struck at 90% of your

purchase price is here more efficient than a standard put of same features, as it also hedges you against a decline but ceases to exist as soon as your stop loss is triggered.

The main advantage of barrier options is that :

3) Their premia are lower than corresponding standard options
   This is intuitive : constraints are put on the payoff, hence it must be that these options are cheaper.

4) Barrier options can offer more transparency than other exotic options
   If an investor has a market view sophisticated enough to qualify for a barrier option or another exotic derivative that also meets his stance, she/he may opt for the barrier option as it will often be more transparent than another suitable derivative.
   E.g., say the investor has the following directional and chronological view : she/he reckons that the spot will very soon rise a lot and then, say a month later, will fall sharply. At least two options may satisfy her/his needs : a lookback put with floating strike of tenor one month (where the strike is the maximum of the underlying over the life of the option), and an U&I put.
   She/he may prefer the U&I put as it is easier to reprice independently and to assess its price ; the lookback option is deemed more exotic.

At this point, one may have noticed that for a D&I put to meet the above conditions and make sense, the barrier has to be set below the strike : indeed, it would otherwise be a disguised standard put, because all scenarios where the standard put is ITM would imply that the corresponding D&I put also is ITM ; now to obtain a cheaper premium there must exist scenarios where a standard put is ITM while a D&I put is not. On the markets, although these products are traded over-the-counter (OTC), the barrier level is usually set around 75% or 80% of the strike.

Payoff profile and sensitivities

Unlike « out » options, the D&I put is straightforward in the sense that its payoff is not subject to opposite effects (like an U&O call or a D&O put which have tensions between getting better moneyness and approaching the deactivating barrier). Its holder therefore doesn’t experience vega negative or theta positive periods, and its delta doesn’t change sign. However, the hedging (especially the delta-hedging) can still be difficult, as close to maturity and slightly above the barrier, the delta is massively negative which implies holding huge quantities of the underlying (possibly more than the notional of the option) and getting ready to sell at a loss large quantities of it as soon as the barrier is breached. Various techniques (such as « barrier shifting ») to mitigate the gap risk exist.

See the graphical representation of the payoff, the delta and the gamma from the aforementioned book by E.Derman :

DOWN-AND-IN PUT:
STRIKE 100, BARRIER 80

PANEL A. VALUE

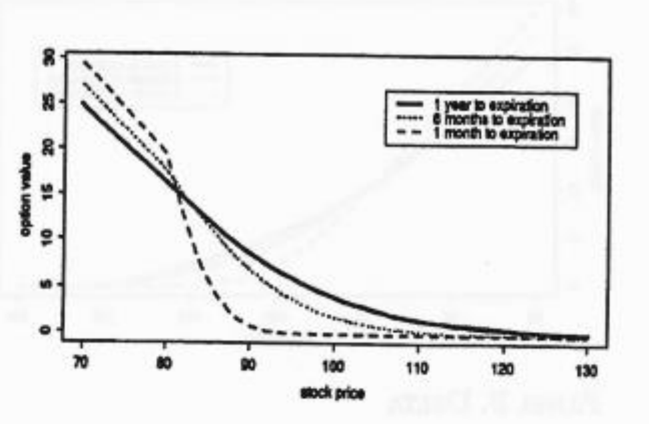

Below barrier, the payoff is that of a standard put struck at 100. The farther (above) from barrier spot is, the lower the probability of knock-in and the premium. At maturity, payoff is a straight line close in shape to the upper dashed line shown between strikes 70 and 80.

PANEL B. DELTA

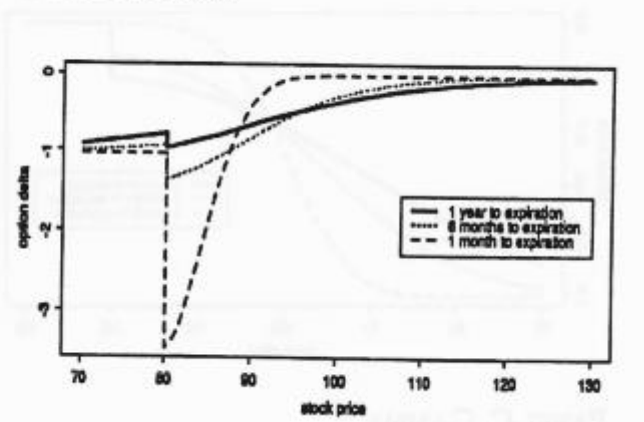

As maturity approaches, delta-hedging (slightly above barrier) is more and more difficult and delta is becoming more and more negative. At maturity and at the barrier, delta is infinite.

PANEL C. GAMMA

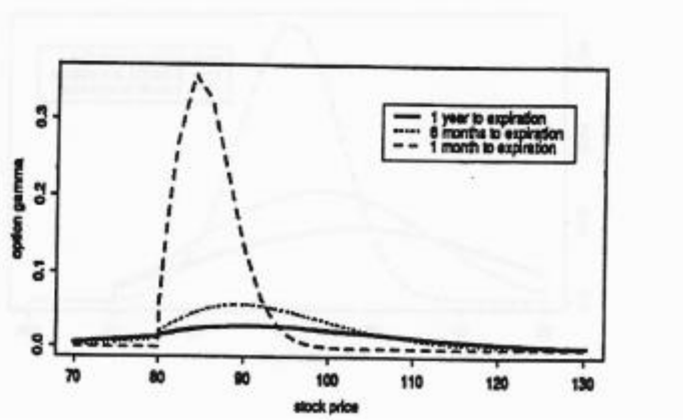

Unlike a standard put, gamma is not highest at the money (ATM). At maturity and at the barrier, it is infinite.

There exists one interesting limit case : imagine that the underlying asset has a poor liquidity, in other words that the price impact of trading it is relatively strong. Second, imagine we live in a world where many traders have the same long D&I delta-hedged put position with big notionals at stake, while their client counterparties wish to take some unhedged exposure by not delta-hedging ; third, these traders delta-hedge very often. Then, in the extreme case, the D&I put will expire worthless almost surely ! (Unless, for instance, the price experiences a large enough downward jump).

Indeed, the traders are delta-negative so their synchronized delta-hedging has an attractive influence on the spot level : whenever spot increases, they sell and when spot falls, they buy. Because the price impact is strong and the notionals hence the amounts of spot traded are huge, the price impact is strong and the spot is virtually pinned to some level close to the spot at inception. Provided the barrier is not too close from this level, and in the absence of jumps, the barrier will remain unhit so the barrier option actually has virtually no value. Graphically, the evolution of the spot would follow some cobweb plot fashion.

Conversely, had the traders been all short the D&I put, their delta-hedging would have had a repulsive effect on the spot, giving it some volatility. The barrier would very likely be breached. So the D&I put would have the same value as a standard put of same features.

II Pricing D&I puts within the standard framework

Implicitly, we work here in the equity world, particularly on single stock options. We will start by pricing the D&I put in the standard, B&S context. This will help us see to what extent this structure appreciates when we switch to a stochastic volatility and jumps framework.

Monte-Carlo simulation

We run 50,000 simulations using the following parameters :

Strike	100
Spot	100
Rebate	0
Time to maturity	1 year
Dividend yield	0%
Constant, annual volatility	20%
Constant, annual continuously compounded risk-free rate	2%
Barrier	80
Frequency of spot level observation	Daily
Corrected barrier	80.49

In order to get a more accurate pricing and to be able to compare our MC pricing to closed-form formulae, which assume continuous monitoring of the spot price, we used the continuity correction given by Broadie, Glasserman and Kou given in « A continuity correction for discrete barrier options ». They provide a correction that accounts for a discrete spot-monitoring instead of a continuous one. For a D&I option, the barrier, say B, becomes B*exp(0.5826*σ*√(T/m)) where m is the number of spot observations and T is the maturity expressed in years.

We thus get in the present scenario a corrected barrier of about 80.49.

Intuitively, any « in » option will get its premium lowered if the spot-monitoring is discrete instead of continuous, as there is less chance of the barrier being hit because we ignore some times where it could be so. Therefore, because we study here in particular a « down » option, the continuity correction consistently gives a higher barrier.

In order to improve our random numbers generation in the MC simulation, we used stratified sampling : we decomposed the [0;1] interval into ten equally-sized sub-intervals (i.e. from [0;0.1] to [0.9;1]) and ensured that inside each of the ten sub-intervals there are ten percent of the random numbers (which are needed to draw normal values).

Lastly, we assume that any day before maturity is a trading day, i.e. that in a year there are 365 (or 366) trading days and not 252, to be able to compare our result to the closed-form formulae once again.

Here are the results obtained :

- Price of the standard vanilla european put : 6.91 (standard error=0.04) (the exact, B&S value is 6.936)
- Price of the D&I put : 5.04 (standard error= 0.04)
- As a consistency check, we also priced the D&O put of same features and found 1.87 (standard error 0.02), which is consistent with the barrier options « parity » relationship : D&O + D&I = Vanilla option.

To get the standard put price, we simulated and averaged 50,000 expiry values using the B&S parameters in the above grid.

To get the D&I and D&O put prices, we simulated 50,000 times the whole spot path which itself contains 365 successive prices, i.e. our MC simulation worked by daily increments, using the well-known equation that governs the evolution of the spot price :

$S_{t+1} = S_t*\exp((r-d-\sigma^2/2)*\Delta t + \sigma*z*\sqrt{\Delta t})$ for all t between inception and maturity, where :

r = constant, annual, continuously compounded risk-free rate

d = continuous, annual dividend yield

σ = constant, annual volatility

Δt = time increment (=1/365 with maturity 1 year and daily increments)

z = a random draw from a standard normal distribution (to simulate the brownian motion)

The prices found are theoretical. In reality, the market quotes for these options would likely be higher due to the downward sloping volatility smile in the equity world, reflecting the fact that tail events, in particular crashes, are actually more likely than what the lognormal distribution used predicts. A D&I put can be seen as an insurance against a market crash.

We could of course estimate some greeks, e.g. the delta, by repricing the option with this time a small increment of the spot, then taking the difference of option prices over the spot increment, with backward, centered or forward scheme.

PDE method (explicit finite differences framework)

In this section, we use the finite differences method, hereby discretizing the partial differential equation followed by the price of the D&I put on a single stock paying no dividend:

$\partial V/\partial t + 0.5*\sigma^2*S^2*\partial^2 V/\partial S^2 + r*S*\partial V/\partial S - r*V = 0$

Where :

- V= option price
- S = underlying stock price
- r = annual risk-free rate
- σ = constant, annual volatility of the underlying stock price

The above PDE is transformed into several differences equations, that are solved in an iterative fashion. Each of its terms is approximated below by finite differences.

Below are the technical details of the process :

We discretize time by creating 250 time periods, and space by creating 200 spot prices from S_{min} to S_{max} (in fact, we will work with log(spot) as this is more convenient for calculations) in a grid. A spot increment is denoted ΔS and a time increment Δt next.

As for boundary conditions, we used the intrinsic value of the option, that is : about 0 for the highest spot value of the grid, at any date, and $K*\exp(-r*(T-t))-min_{spot}$ for the smallest spot value of the grid. This seems a good enough approximation as deep ITM or deep OTM, the time value of an option is very small.

In this method, we model the process of the underlying price S (in the risk-neutral world) by a standard process which will become in the second part of this study a jump and stochastic volatility process. We then apply Ito's lemma to the option function to get the PDE.

Next, we need to approximate the terms in the PDE presented above (for all points strictly inside the grid):

- dV/dS (the delta of our option) is approximated using the centered scheme (for its nicer mathematical properties compared to the backward and forward scheme ; the centered scheme is the average of the forward and backward schemes).
  Let us denote $f_{i,j}$ the value of our D&I put at date i and for a spot level j, then :
  $dV/dS = (f_{i,j+1} - f_{i,j-1}) / 2\Delta S$

- d²V/dS² (the gamma of the option) is approximated as the derivative of the delta, therefore we compute the difference between the forward scheme and the backward scheme:
  $d^2V/dS^2 = ((f_{i,j+1} - f_{i,j}) / \Delta S - (f_{i,j} - f_{i,j-1}) / \Delta S)/ \Delta S$

- dV/dt (the theta of the option) is approximated using the forward scheme, so as to get a direct link between the option premium at two consecutive dates :
  $dV/dt = (f_{i+1,j} - f_{i,j}) / \Delta t$

Because the underlying spot price follows a geometric brownian motion, we now change variable, and choose Z=log(S) as calculations are then simpler.

We replace the three above terms in the PDE and obtain this relationship :

for all $f_{i;j}$ strictly inside the grid (i.e. not on a boundary):

$$f_{i;j} = \Psi_+ * f_{i+\Delta t ;j+\Delta\log(S)} + \Psi_0 * f_{i+\Delta t ;j} + \Psi_- * f_{i+\Delta t ;j-\Delta\log(S)}$$

where :

$$\begin{cases} \Psi_+ = (\Delta t/(2\Delta\log(S)))*(r-\sigma^2/2) + (\Delta t/(2\Delta(\log(S)^2)))*\sigma^2 \\ \Psi_0 = 1 - r*\Delta t - \sigma^2*\Delta t/\Delta(\log(S))^2 \\ \Psi_- = - (\Delta t/(2\Delta\log(S)))*(r-\sigma^2/2) + (\Delta t/(2\Delta(\log(S)^2)))*\sigma^2 \end{cases}$$

This relationship is completed by the known values on the three boundaries, i.e when i= maturity and/or j= $\log(S_{min})$ or j= $\log(S_{max})$.

We can now solve backward in the grid (i.e. we start by using the prices at maturity, which are known) for the price of the D&I put at inception.

To ensure stability of this numerical analysis scheme, we chose the spot increment $\Delta\log(S)$ such that the three multiplicative coefficients Ψ_+, Ψ_0, and Ψ_- remain positive (and they sum up to $1 - r*\Delta t$ as of Lax’s theorem) and such that there is a proportional relationship between

$\Delta log(S)$ and $\sqrt{\Delta t}$ for matters of convergence towards the right price. These coefficients are like the risk-neutral probabilities of states « up », « middle » and « down » in a trinomial tree.

A key element of this scheme, that has a direct impact on the price obtained, is the size of $\Delta log(S)$ (i.e. the value between two consecutive spot prices). We here chose $\Delta log(S) = 0.022$ and $\Delta t = 0.004$ (indeed we have 250 periods in our grid, as said at the beginning of this section). Theoretically, the right price is obtained when both $\Delta log(S)$ and Δt converge towards 0.

The price found, using the same parameters as in the previous section, i.e. without using jumps or stochastic volatility is 5.17.

The difference found with other pricing methods lies in several factors :

- The relatively long maturity : it increases the errors made by the derivatives approximations thereby reducing the accuracy of the model

- The spot increment chosen : it needs to be such that the multiplicative (Ψ) coefficients remain positive and between 0 and 1 (and ideally roughly equal to 1/3 each) but also such that the extreme spot values in the grid are sensible with regard to the maturity chosen, i.e. deltax must be basically proportional to the maturity chosen. However, the finite differences scheme can be quite sensitive to deltax and does not necessarily increase the option price linearly with deltax (increasing deltax can sometimes reduce the option price). To sum up, the choice of the spot increment in the grid is crucial and can be a source of imprecision in pricing. A usual choice is $\Delta log(S) = \sigma * \sqrt{(3\Delta t)}$ which we use here.

This is why in order to test the convergence between the abovementioned Monte-Carlo method and the PDE method, we reduced the maturity to 15 days and found that with a volatility of 20%, the $P_{D\&I}$ is virtually worthless with both methods. This doesn't come as a surprise, as a volatility of 20% means the underlying stock price is « quiet », and when given only 15 days it virtually never hits the barrier, due to lack of time to do so and to lack of volatility. Actually, the $P_{D\&I}$ value starts exceeding one cent when the volatility level is 33% or more, *ceteris paribus*, i.e. for any volatility below 33% its value is less than a cent.

With a volatility of 50%, the Monte-Carlo method (with 50 000 simulations) and the PDE method (deltax=0.02) show price discrepancies of a cent, clearly showing the convergence, as expected.

Image method

We first present this method, based on the findings in the paper by P.Buchen « Pricing european barrier options ».

In that paper, Buchen reveals a novel way of pricing barrier options, which appears to be simpler than the then usual expectations methods or than the study of the law of the min/max of the underlying asset. The latter are mathematically-intensive, for example they require the determination of the risk-neutral density of the underlying asset price (in « Prices of state-contingent claims implicit in option prices », 1978, Breeden&Litzenberger approximate the density by deriving twice the option premium relative to strike, then by assuming smoothness of the density around the strike and using a butterfly spread, lastly by assuming that the density is a step function, increasing then decreasing, in Riemann's integrals style. This gives a clue as to the quantitative background required). In 1995, Ritchken had performed numerical methods of barrier option pricing using binomial and trinomial lattices, experiencing the drawback that the barrier level can fall between two branches of the tree, thus not properly accounting for the dynamics of the underlying asset.

By contrast, Buchen's method requires a fairly limited amount of mathematics :

Buchen sorts the eight barrier options according to their active domain (the interval where the barrier has not been hit yet, e.g. [barrier ;+∞] for a down option), their expiry condition (i.e. what they are worth at expiry if the barrier was not hit, either a plain vanilla option, or a rebate, or nothing if no rebate) and their boundary condition (the value of the barrier option once the barrier is breached). He also decomposes vanilla options in their main components, i.e. their elementary constituents, which are the four elementary solutions of the B&S partial differential equation (PDE) presented in the previous section.

These constituents are, for a B&S call, $S(t)*\phi(d_1)$ and $\exp(-r(T-t))*\phi(d_2)$, and for a B&S put $S(t)*\phi(-d_1)$ and $\exp(-r(T-t))*\phi(-d_2)$, where :

$d_1=(\log(S(t)/K)+(r+0.5*\sigma^2)*(T-t))/(\sigma*\sqrt{T-t})$

$d_2=d_1-\sigma*\sqrt{T-t}$

ϕ is the standard normal cumulative function

In this fashion, we get a high level of granularity and have an analytical view of the very components of options, and of the discrepancies between barrier options.

Once this has been done, we have the groundwork to welcome the key contribution of the paper : Buchen defines, for each of the four elementary solutions of B&S PDE, their *image solutions*, which are graphically their symmetric in the log-space (that is, when considering the log of the underlying price) relative to the (vertical line formed by the) barrier, i.e. the straight line of equation x=barrier.

Thus, Buchen rightly uses the put-call symmetry relationship :

Definition 1 *Let $u(x,t)$ be any solution of the BS-pde. Then the image of $u(x,t)$ relative to the barrier $x = b$ is defined to be the function*

$$u^*(x,t) = \left(\frac{b}{x}\right)^{\alpha} u\left(\frac{b^2}{x}, t\right).$$

Where b is the barrier, x the spot, t time to maturity (which is usually denoted rather T-t), and $\alpha = 2r/\sigma^2 - 1$ where r is the risk-free interest rate and σ is the underlying asset price volatility. The first term is a corrective term aimed at accounting for the drift, which is not symmetrical.

The symmetry is in log space in the above formula, consequently we find indeed that $\log(b^2/x) = 2\log(b) - \log(x)$ is the symmetric of log(x) relative to the log of the barrier, log(b).

These image solutions satisfy all the relations needed relative to the B&S operator.

All barrier options can now be expressed in terms of the eight fundamental solutions (four plus their images). This means the volatility smile only can be used to price barrier options : by virtue of this method, one transforms the american feature of the barrier into an european-only combination of options (assuming one is able to replicate the symmetric of a standard option).

Also, it should be noted that thanks to parity relationships laid out by Buchen, the knowledge of any one of barrier options price gives knowledge of all other barrier option prices.

Below is the premium of the D&I put expressed in terms of standard puts, themselves being linear combinations of two of the fundamental components described above.

$P_{D\&I} = P_K^* + P_b - P_b^*$ on its active domain i.e. for S(t) in $]b ;+\infty]$

Where : P_y is a standard put struck at y

P_y^* is the symmetric of a standard put struck at y (in log space)

Using the same parameters as in the two previous sections (Monte-Carlo and PDE), we find a price of the D&I put of 5.096, broken down as follows :

$P_{D\&I}(100;1\ \text{year}) = 34.113 + 3.552 - 32.569$

The prices of the three put components are computed using the well-known B&S closed-form formula :

$P(t) = -S(t)*\phi(-d_1) + K*\exp(-r(T-t))*\phi(-d_2)$ for any t between 0 and T

Where :

$d_1=(\log(S(t)/K)+(r+0.5*\sigma^2)*(T-t))/(\sigma*\sqrt{T-t})$

$d_2=d_1-\sigma*\sqrt{T-t}$

ϕ is the standard normal cumulative function

Closed-form formula

Let us check now how precise the three pricing methods used above are, by comparing the prices obtained to the price given by the closed-form formula. The latter is taken from the second part of « The ins and outs of barrier options » by E.Derman.

Derman provides the formula for a D&O put that he derives from the known similarities with the U&O call. From this D&O price we infer the price of the D&I put of same parameters by using the well-known relationship : $P_{standard} = P_{D\&O} + P_{D\&I}$

$P_{D\&I} = P_{standard} - K*\exp(-r(T-t))*[\Phi(-d_1+\sigma*\surd(T-t)) - \Phi(-X_2+\sigma*\surd(T-t)) + \Phi(-Y_1+\sigma*\surd(T-t)) - \Phi(-Y_2+\sigma*\surd(T-t))] - S(t)*[\Phi(-d_1) - \Phi(-X_2) - \Phi(-Y_1) + \Phi(-Y_2)]$

Where :
- d_1 is the usual B&S component, defined above
- $X_2 = d_1$ replacing K with b, the barrier
- $Y_1 = d_1$ replacing S/K in the log with $b^2/(S*K)$
- $Y_2 = d_1$ replacing S/K in the log with b/S

The formula is thus simplified here as we are in the special case where the continuous dividend yield, the power alpha and the rebate are all equal to zero.

Note that we could have used, alternatively, Derman's observation that a D&O call (resp. an U&O put) is almost a call (resp. put) spread : e.g. for the call spread, the short call is held in a proportion equal to the corrective term introduced earlier, and its strike is equal to $K*S^2/b^2$.

By replacing the parameters chosen (see in the MC subsection) direclty into the above formula, we find a price of 5.096.

The Monte-Carlo method (resp. the PDE method) therefore tend to under(resp. over) estimate slightly the option price, for the reasons we discussed previously, while as expected the image method leads to the exact price in a symmetric world.

Let us see whether this still holds in an asymmetric context.

III Pricing D&I puts within an asymmetric environment

This part is motivated by the deficiencies of standard frameworks like Black&Scholes, which hypotheses are too simplistic. This simplicity is obviously a weakness, but also a strength (everyone can understand it), that is why B&S is still so widespread in banks.The B&S model makes strong assumptions that are not reasonable, such as :

- the ability to delta-hedge continuously and without any transaction costs
- the underlying is not illiquid
- any day is a trading day (this triggers problems, mainly regarding the assessment of volatility)
- investors are rational
- there are no arbitrage opportunities
- one can short-sell a stock, and for any amount
- the yield curve is flat
- the underlying stock doesn't distribute any dividend over the life of the option
- the underlying stock price follows a lognormal distribution (this usually underestimates the left tail)
- the volatility of the underlying stock price is constant through time
- the underlying stock doesn't experience any jumps, including default

In the following part, we focus on the last two assumptions which appear to be critical : obviously, according to the political, economic and financial context through the life on an option, its underlying stock volatility changes (potentially a lot). Moreover, jumps are a recurrent phenomenon on the stock market (it is probably less so in the FX world, at least for major currencies, as this market is much deeper and more liquid). For instance, one case of jump frequently observed is a downward jump just after a distribution of dividends, of the amount the dividend per share.

In what follows, we create an environment closer to the market by adding jumps (positive, negative, and to default) to the underlying price dynamics as well as by implementing a stochastic volatility framework. We thus also build transition jumps (i.e. jumps that trigger another volatility regime).

We will use the fact that some of the methods employed do not rely on symmetry arguments (namely, the Monte-Carlo simulation, the closed-form formula and the finite differences method) while the image method in essence depends on it. We should therefore be able, by comparing the prices obtained in this asymmetric environment, to assess the extent of the discrepancies, and in turn detect potential mispricing and hedging risks.

Of course, we simplify a lot reality. For sake of simplicity, we work here on a two-regime model, and again on single stock options. So, we won't be dealing with correlation here. In addition, whenever a jump occurs in our model, its size is known and constant. We could complexify the model by creating a distribution of jumps that would randomly determine the size and/or the sign of jumps. Moreover, we assume that at any time over the life of the option, the volatility of the underlying price assumes one of only two predetermined values, basically corresponding to a high (resp. low) volatility regime. The jump-to-default intensity is set higher in the high volatility regime than in the low volatility regime, which fits best reality in most situations. We also assume a flat yield curve throughout, whatever the regime. More generally, we keep the B&S assumptions mentioned above other than those regarding volatility and jumps, plus note that we assume a null rebate and we continue to perform a continuity correction for the daily spot observations instead of continuous monitoring.

Monte-Carlo simulation

We run 100,000 simulations using the following parameters :

Strike	100
Spot	100
Rebate	0
Time to maturity	1 year
Dividend yield	0%
Volatility of regime 1 (σ_1)	20%
Volatility of regime 2 (σ_2)	50%
Risk-free rate	2%
Barrier	80
Intensity of positive jump in regime 1	2
Intensity of negative jump in regime 1	1
Intensity of positive jump in regime 2	3
Intensity of negative jump in regime 2	5
Intensity of switching regimes 1->2 (λ_2)	5
Intensity of switching regimes 2->1 (λ_1)	4
Initial regime	Regime 1
Frequency of spot level observations	Daily
Corrected barrier	80.9
Default intensity in regime 1	0.01
Default intensity in regime 1	0.03
Size of positive jump in regime 1	5%
Size of negative jump in regime 1	-4%
Size of positive jump in regime 2	5%
Size of negative jump in regime 2	-8%

Like in the previous MC simulation (in the B&S framework), we will use a continuity correction in order to get a more accurate pricing. This time however, we can no longer use Broadie, Glasserman and Kou's formula as we now have stochastic volatility instead of a constant volatility, yet the formula uses a single, constant level of volatility.

We will therefore use the following approximation : we will compute a weighted average volatility by using the two volatility levels and their respective transition intensities. We then have : $\sigma_{average} = (\lambda_2/((\lambda_2+\lambda_1))*\sigma_2 + \lambda_1/((\lambda_2+\lambda_1))*\sigma_1) = 5*50\%/9 + 4*20\%/9 \approx 37\%$.

We are aware of the flaws of the above calculation. However the transition intensities, very close to each other, should help mitigate the error made by this approximation. Moreover, this continuity correction shifts the barrier by a small, almost negligible amount anyway. We now get a corrected barrier of : 80*exp(0.5826*37%*sqrt(1/365)) ≈ 80.9 instead of 80.

Consistently, we obtain as explained in the previous Monte Carlo section, a higher barrier.

Note that the size of jumps is defined relative to the spot price on the previous day.

Note also that contrary to intuitive belief, defining one regime-switching intensity does not define the other one : one could think that switching n times say from 1 to 2 (and 1 being the initial regime) implies to switch either (n-1) or n times from 2 to 1. In fact, the intensity should be interpreted rather in terms of the « weight » of each regime, that influences the overall volatility.

However, the above remark says nothing regarding the time spent respectively in each volatility regime : it could be in theory that the spot switches the same number of times between regimes, but however remains in one of the two regimes much longer than in the other one, which would considerably affect its behaviour : it would be almost as if there was only one regime (the one in which the spot stays longer). But by construction (remember the Poisson process used to count jumps), the transition times should be on average uniformly spread over the life of the option.

In the above grid of parameters, we didn't choose the jumps sizes totally arbitrarily. First, in order to be realistic we set the negative jump size in the high volatility regime higher (in absolute value) than that of the low volatility regime. Second, we chose these jumps sizes such that they correspond to jumps sizes « available » in the grid of the PDE method (see next section) : indeed, in the grid any jump falls exactly on a node, by assumption, so in order to compare the models we need to choose roughly the same jumps sizes. In other words, we first looked in the PDE method grid the size of the jumps (which is directly linked to the spot increment), then set it for the Monte-Carlo simulation.

In order to improve our random numbers generation, we decomposed the [0;1] interval into ten equally-sized sub-intervals (i.e. from [0;0.1] to [0.9;1]) and ensured that inside each of the ten sub-intervals there are ten percent of the random numbers (which are needed to draw normal values), conform to stratified sampling.

We kept using the daily-incremental formula for the evolution of the spot price presented in the previous MC section, except this time we used contingency conditions on the volatility and we added the prospective jumps.

Regarding the jumps, of course we made sure that in the (rare) occurrence of a jump-to-default, there is not in addition a negative jump (although this « double jump event » is highly unlikely), so that the spot price is indeed bounded below by zero i.e. cannot become negative and also because this would not make sense as in a way the jump-to-default embeds the negative jump.

- Price of the standard vanilla european put : 19.01 (standard error = 0.07)
- Price of the D&I put : 18.55 (standard error = 0.07)
- As a consistency check, we also priced the D&O put of same features and found 0.47 (standard error = 0.01)

The D&O put is virtually worthless in this context because the negative jumps and the jumps-to-default together with a volatility of 50% in the high regime (and an average volatility that almost doubled compared to the MC simulation in the B&S context) almost ensure that the barrier will be breached during the life of the option.

Equivalently, we notice the massive appreciation of the D&I put in this new context (it was worth roughly 5 in the previous, B&S section). Thus, in models that integrate stochastic volatility and jumps, the traditional rationale for buying barrier options (namely, their low cost) is often mitigated for « In » options as their price converges towards the corresponding standard option. Indeed, these new features make a « barrier event » much more likely.

PDE method (explicit finite differences framework)

In this section, we use the finite differences method, hereby discretizing the partial integro-differential equation followed by the price of the D&I put.

Below are the technical details :

We start from the same framework as in the symmetrical context. The boundary conditions are : 0 for a spot equal to S_{max} , the value of the standard put of same features at S_{min}, (as S_{min} is below the barrier) and either 0 or the value of the standard put of same features at maturity, according to the spot level relative to the barrier.

We now use two differents regimes of (brownian) volatility, each one corresponding to one grid. Jumps within and between regimes are described by Poisson processes which intensities (i.e., average number of occurrences over all the time periods) are arbitrarily chosen. These intensities are chosen with common sense though, e.g. negative jumps intensities are set higher than positive jumps intensities in order to reflect the alleged true distribution of the underlying price that has a fatter left tail and a thinner right tail relative to a lognormal distribution ; and obviously negative jumps intensities are much higher than jump-to-default intensities) . There can be positive, negative and default jumps inside each of the two regimes, plus transition jumps (i.e. from one regime to the other), which all contribute to the total volatility being higher than the brownian volatility. For sake of simplicity, we assume as stated previously that any jump size is a multiple of the spot increment $\Delta\log(S)$, that is to say that any jump falls exactly on a node on one of the two grids. Any jump that would fall outside of the grids is valued using the intrinsic value of the option given the spot price right after the jump occurred.

In order to have a finer analysis, we extended the grid to 100 time periods and 200 spot levels. This helps to have more precision in the pricing as a common problem in tree-like approaches to barrier option pricing is that the barrier can fall between two nodes.

We now have to include jumps in our PDE. This makes it a PIDE (partial integro-differential equation) :

$$dV/dt - r*V + (r-\Sigma\lambda_i y_i)*V*dV/dS + (\sigma^*)^2/2*V^2*d^2V/dV^2 + (\Sigma\lambda_i*\Delta V_i) = 0$$

where ΔV_i measures the change in option price subsequent to the jump n°i, and σ^* is the volatility of either of the regimes used.

In our study we use three jumps : positive, negative and default.

Accounting for the jumps we obtain this relationship :

for all $f_{i;j}$ strictly inside the grid (i.e. not on a boundary):

$$f_{i;j} = \Psi_+ * f_{i+\Delta t;j+\Delta\log(S)} + \Psi_0 * f_{i+\Delta t;j} + \Psi_- * f_{i+\Delta t;j-\Delta\log(S)} + \Sigma(\Psi_{\Delta k} * f_{i+\Delta t;j+\Delta k})$$

where :

$$\begin{cases} \Psi_+ = (\Delta t/(2\Delta\log(S)))*(r-\sigma^2/2) + (\Delta t/(2\Delta(\log(S)^2)))*\sigma^2 \\ \Psi_0 = 1 - r*\Delta t - \sigma^2*\Delta t/\Delta(\log(S))^2 \\ \Psi_- = -(\Delta t/(2\Delta\log(S)))*(r-\sigma^2/2) + (\Delta t/(2\Delta(\log(S)^2)))*\sigma^2 \\ \Psi_{\Delta k} = \lambda_k * \Delta t \end{cases}$$

For matters of comparison, we use of course the same intensities and jump sizes as those shown in the previous section (see grid in « Monte-Carlo simulation within an assymetric environment »).

In this context, we find a much higher price than under standard, B&S conditions : about 17.4.

This is explained by the very volatile regime 2 together with a relatively high transition intensity from 1 to 2.

The discrepancy in pricing comes from several biases, some of which were already mentioned in the PDE section of the first part. In this part, we add even more error due to the jumps feature of this new context : the PDE method imposes a rigid structure : the grid has fixed increments (although some methods exist to vary spot increments according to spot level) and therefore we are bound to work with approximations when it comes to the jump size : in other words, we fixed the jumps sizes for the Monte-Carlo simulations which is fine for Monte-Carlo, but in our grid, we cannot jump from exactly these percentages. Indeed, we assume that any jump falls exactly on a node of the grid, not between two nodes. Besides, the explicit scheme used is known to be less effective than the implicit scheme and especially than the Crank-nicholson scheme, which is basically a mix of the two.

In order to check the convergence between MC and PDE methods, we thus chose jumps sizes that approximately match spot increments in the grid and used these values in our MC simulation.

Image method

Remember, we used the below relationship :

$P_{D\&I} = P_K^* + P_b - P_b^*$ on its active domain i.e. for S(t) in]b ; +∞]

Where : P_y is a standard put struck at y

P_y^* is the symmetric of a standard put struck at y (in log space)

In the new context of assymetry of the volatility smile, we will price again the D&I put by pricing each of its three components, with the finite differences method, and will check the size of the discrepancy.

First, we price directly the D&I put with the same finite differences and MC methods and parameters as the one used previously, and find a price of ≈ 17.4 with PDE and ≈ 18.55 with MC (this was done in the previous section).

Then, we price separately Buchen's three constituents of the D&I put. We use the payoffs at maturity for the simulations, and the parameters used in the two previous sections. The payoffs at maturity are : $\max(0 ;K-b^2/S) + \max(0 ;b-S) - \max(0 ;b-b^2/S)$, using both the formula framed in blue above, and the definition of the symmetric of an option in the image method (equation page 12). In this MC simulation, we therefore use the closed-form formulae given by Buchen. In order to adapt them to our stochastic volatility and jumps context, we simulated the whole path. In other words, the formula above is evaluated 100,000 times at the end of 365 observations during which volatility can change regime, and jumps can occur.

Pricing separately the three components of the D&I put with 100,000 MC simulations yields a price of ≈ 22.55.

We find therefore a material discrepancy here in this particular context, as Buchen's formula leads to overestimate the D&I structure. This tends to restrict Buchen's analysis validity to a symmetric environment.

In order to test again Buchen's argument validity, we generate a smile that is less assymetric, namely with smaller and less frequent jumps and relatively low volatilities.

Strike	100
Spot	100
Rebate	0
Time to maturity	1 year
Dividend yield	0%
Volatility of regime 1 (σ_1)	20%
Volatility of regime 2 (σ_2)	40%
Risk-free rate	2%
Barrier	80
Intensity of positive jump in regime 1	1
Intensity of negative jump in regime 1	2
Intensity of positive jump in regime 2	3
Intensity of negative jump in regime 2	4
Intensity of switching regimes 1->2 (λ_2)	5
Intensity of switching regimes 2->1 (λ_1)	2
Initial regime	Regime 1
Frequency of spot level observations	Daily
Corrected barrier	80.9
Default intensity in regime 1	0.01
Default intensity in regime 1	0.02
Size of positive jump in regime 1	3%
Size of negative jump in regime 1	-4%
Size of positive jump in regime 2	3%
Size of negative jump in regime 2	-8%

We now find that the price of the D&I put priced directly is ≈ 17.88 (standard error 0.08) ; as expected the price is still high compared to a B&S framework due to the negative jumps and stochastic volatility, but lower than in the previous simulation. The price of the D&I obtained from its components (image method) is ≈ 20.22. Thus once again, the image method overestimates the D&I put, with the parameters used. Note that the sensitivity of the option price to a reduction in the volatilities and jumps used in the simulations looks fairly low.

Let's test a scenario even closer to B&S. In particular, we use the following parameters :

Strike	100
Spot	100
Rebate	0
Time to maturity	1 year
Dividend yield	0%
Volatility of regime 1 (σ_1)	20%

Volatility of regime 2 (σ_2)	30%
Risk-free rate	2%
Barrier	80
Intensity of positive jump in regime 1	1
Intensity of negative jump in regime 1	1
Intensity of positive jump in regime 2	2
Intensity of negative jump in regime 2	2
Intensity of switching regimes 1->2 (λ_2)	2
Intensity of switching regimes 2->1 (λ_1)	2
Initial regime	Regime 1
Frequency of spot level observations	Daily
Corrected barrier	80.9
Default intensity in regime 1	0.001
Default intensity in regime 1	0.01
Size of positive jump in regime 1	3%
Size of negative jump in regime 1	-3%
Size of positive jump in regime 2	3%
Size of negative jump in regime 2	-6%

We now find that the price of the D&I put priced directly is about 8.5 (standard error 0.06) while the price of the D&I obtained from its components (image method) is about 18.9. This time, while the direct pricing leads to a price that dramatically decreased, the pricing using the three constituents has decreased only a little.

Again, the image method doesn't hold in the case studied, which is rather close to a symmetric volatility smile, though.

IV Conclusions

It seems that the image method to price barrier options is valid only in the specific case where the volatility smile is symmetrical.

Indeed, while the D&I put price perfectly matches the exact price (the one found with the closed-form formula) within a symmetrical framework, there seems to have a sizeable discrepancy when it comes to an asymmetrical volatility smile : when we add jumps and stochastic volatility to the standard framework, the prices start to differ by material amounts, probably too much for a bank looking for accurate pricing.

In particular, it seems that the image method repeatedly gives a relatively high price with little sensitivity to the level of the jumps and volatilities used ; in other words when using a very asymmetrical smile (large and frequent jumps) the pricing ought to be relatively accurate, however when simulating a smile close to being symmetrical, the discrepancy is largest.

V Bibliography

- « A continuity correction for discrete barrier options », Broadie, Glasserman and Kou (1997), Mathematical finance, Vol 7, No 4, 325-348
- « Pricing European barrier options », Buchen (2006)
- « The ins and outs of barrier options », 1 and 2, Derman (1996/97)
- « On pricing barrier options », Ritchken (1995), The journal of derivatives, 19-28
- « Breaking down the barriers », Rubinstein and Reiner (1991), Risk, 28-35
- « Hedging complex barrier options », Carr and Chou (2002), 6-8
- « Prices of state-contingent claims implicit in option prices » (1978), Breeden and Litzenberger

VI Thanks

Thanks to Mr Henrotte for his support and advice.