

A GENERIC APPROACH TO MANAGE GAP RISK

Frank Koster
Private communication
frank.koster@gmx.de

Daniel Oeltz
RIVACON GmbH
d.oeltz@rivacon.com

Abstract

We present a generic framework to manage gap risk for financial products involving large digital risks due to discontinuous payouts or barriers for products priced by solving a partial differential equation (PDE). The approach modifies the boundary conditions of the pricing PDE so that a given delta limit is not breached. In comparison to other techniques such as simply shifting the barriers it produces much smaller price offsets to the original fair value. We also discuss how the approach could be applied to Monte Carlo pricing.

1 Introduction

Most models used in the financial industry to price derivatives assume that the underlying can be liquidly traded in continuous time. For derivatives with discontinuous payoff like barrier options a violation of this model assumption introduces a substantial gap risk. Consider a short position in a continuously monitored Down and Out Put (DOP) option. In a continuous time model the delta may become very large if the spot is close to the barrier shortly before maturity. If spot jumps down breaching the barrier the product expires worthless, but the trader still has a large long position of shares which he can not immediately unwind. This may lead to either large gains (if spot moves up again) or losses (if the spot goes further down) even in common market situations. In critical market situations one often observes that down swings are clustered so that there may be a lot of barrier hits on different underlyings at the same time. In such situations severe losses are the very likely consequence.

A naive approach to circumvent the effect of a large delta is to simply cut off the delta by a given delta limit.

There are mainly two disadvantages for this method. First, just limiting the delta in this way it is not clear what the impact is on the hedge performance, and hedging costs may increase, especially in situations where the spot is close to the barrier but the barrier is not breached. This is also reflected in the fact that the price is unaffected by this method, which makes it also difficult to decide how much should be provisioned for the gap risk. Second this approach is not able to compensate losses which might occur if markets go down for all underlyings of a certain sector where a lot of barrier hits may occur and losses may accumulate.

For these reasons banks apply more sophisticated techniques called management techniques or overhedging to account for the gap risk.

The basic principle of these techniques is to find a so-called managed product which is conservative in the sense that for all market parameters the price is higher (for short positions) or lower (for long positions) than the price of the original product and which has a smaller digital risk in the sense of smaller maximum absolute deltas.

If the original product is sold to a client the resulting position is booked with the managed product into the risk system. Since it is unlikely that we can fully charge the difference between fair and managed price from the customer, the profit of the trade is reduced. Thus we need to find a managed product with optimal trade-off between management margin and reduction of the digital risk.

Let us revisit the example with the DOP from the beginning. One possibility to get a managed product for the DOP is to simply shift the barrier down. Obviously this reduces the barrier hit probability and the managed product will become more

expensive. At the same time, if the spot is close to the original barrier, the delta for the managed product will in general be much smaller (in absolute terms) reducing the delta risk. When the unshifted barrier is hit, the original product knocks out and the trader unwinds his hedge by selling the shares he is long. Here, the trader benefits from two things: First, he faces a (much) smaller gap risk as he has to sell less shares at the lower price if the spot price moved further down since the delta of the managed product is smaller than the delta of the unmanaged product. Second, the trader realizes a windfall profit as the managed product still has a positive value when it is booked out which compensates losses accounting to the gap risk.

This classical approach of shifting the barrier may be difficult or even impossible to apply to more complex products due to the side constraint that the managed product must be conservative. In addition, simply shifting the barrier is not optimal with respect to the fore mentioned trade-off between managing gap risk and the additional margin.

To circumvent these problems we propose a framework which guarantees that the resulting price is conservative in all situations and which can be generically applied to all kinds of products which can be priced by backward solving a parabolic equation. Moreover it may be much cheaper than the classical managed product resulting in lower hedging costs.

For the ease of presentation we restrict ourselves to the case of short positions in the following. The results also directly carry over to long positions.

To control the digital risk we introduce the parameter δ_{max} which must be chosen a-priori according to the liquidity of the underlying (maybe estimated from numbers such as market capitalization or the average traded volume) and the notional of the trade.

The basic idea of our approach is to compute the price $\tilde{V}(S, t)$ by backward solving the usual pricing PDE with the additional constraints that the delta nowhere exceeds δ_{max}

$$|\partial_S \tilde{V}(S, t)| \leq \delta_{max} \text{ for all } S, t. \quad (1)$$

Here, we omitted other market parameters such as volatilities, interest rates and dividends for the ease of notation. The second constraint, that the managed product price must be larger or equal the unmanaged price under all conditions, can be written as

$$\tilde{V}(S, t) \geq V(S, t) \text{ for all } S, t. \quad (2)$$

where $V(S, t)$ is the solution of the PDE for the unmanaged product.

2 The algorithm

We assume that a finite difference scheme is used to solve the backward PDE for the price $V(S, t)$ on a spot grid s_i , $0 \leq i \leq N$ and a time grid t_j , $0 \leq j \leq M$. The price values on the grid are $v_i^j = \{V(s_i, t_j)\}$ and the vector of all values for a given time slice t_j is denoted by v^j . The final payoff is given by a possibly discontinuous function f and for the ease of notation we assume that there is at most a continuous lower barrier at s_0 with possibly time-dependent but continuous Dirichlet boundary values $b(t)$. This setting also covers discrete barriers as the respective pricing PDE is solved by subsequently solving PDEs similar to that with typically discontinuous final payoffs and no explicit conditions at one or both domain boundaries. If the product converts at a barrier into another barrier product the second product also needs to be computed with barrier management to get the right boundary values b . Using this notation the algorithm is.

Algorithm BM

- set $\tilde{v}_0^M = b(T)$ and $\tilde{v}_i^M = f(s_i)$ for $i > 0$
- compute the smoothed final payoff

$$v^M := \text{Smooth}(\tilde{v}, \delta_{max})$$
 where *Smooth* is defined at the end of this section.
- for $j=M-1$ down to 0
 - solve for v^j from v^{j+1} using the existing PDE solver but with a modified boundary condition at s_0 : If

$$v_0^{j+1} > b(t_{j+1}) \quad (3)$$

use Neumann boundary conditions

$$\partial_S V(t_j, s_0) = \delta_{max}$$

and regular Dirichlet conditions

$$v_0^j = b(t_j)$$

otherwise.

To explain the rationale using Neumann conditions we go backwards from t^M to t^0 . Condition $v_0^M > b(T)$ indicates that there is an incompatibility between the smoothed final payoff and the boundary values. Therefore using the simple

Dirichlet condition for time slice $M - 1$ would most likely cause a singularity in the pricing PDE and thus huge deltas in v^{M-1} which may violate our delta restriction. In addition, numerical issues may be introduced due to the singular behaviour.¹ To avoid violations of the delta constraint in time slice $M - 1$ we might solve an optimization problem which tries to find the smallest boundary value $b^* \geq b(t^{M-1})$ such that solving one step of the pricing PDE with Dirichlet values b^* would still satisfy the delta constraint. Because v^{M-1} depends monotone on b^* the minimum is achieved at either $b^* = b(t^{M-1})$ or if $\partial_S V(s_0) = \delta_{max}$ as also $\partial_S V$ satisfies a maximum principle. This argument repeats for all other time steps.

For $j < M - 1$ it is possible to use a more sophisticated rule to switch between Neumann and Dirichlet conditions. One first computes a proxy \tilde{v}_0^j for v_0^j , for example, by linear extrapolation and uses Neumann conditions if $\tilde{v}_0^j > b(t_j)$.

We now describe the algorithm smooth to compute the smallest vector $v \geq u$ which fulfills the delta constraint.

Algorithm $v := Smooth(u, \delta)$
// right sweep
 $w_0 := u_0$
For $i = 1$ to N
 $w_i = \max(u_i, w_{i-1} - \delta(s_i - s_{i-1}))$
// left sweep
 $v_N := w_N$
For $i = N - 1$ to 0
 $v_i = \max(w_i, v_{i+1} - \delta(s_{i+1} - s_i))$
end

Proposition Vector $v := Smooth(u)$ is the solution of the optimization problem

$$\min_w \|w - u\| \quad \text{s.t. } w \geq u \quad \text{and} \quad \left| \frac{w_{l+1} - w_l}{s_{l+1} - s_l} \right| \leq \delta,$$

for an arbitrary vector norm $\|\cdot\|$.

It is easy to see that the right sweep finds a vector $w \geq u$ such that $w_i - w_{i-1} \geq -\delta(s_i - s_{i-1})$. It is also the smallest vector with these properties. Assume there is $w' \geq u$ with $w'_i - w'_{i-1} \geq -\delta(s_i - s_{i-1})$, but $Q := \{i : w'_i < w_i\} \neq \emptyset$. Because $w_0 = u_0$, $i_0 := \min(Q) > 0$. We know that $w_{i_0} = w_{i_0-1} - \delta(s_{i_0} - s_{i_0-1})$ as otherwise

¹The case $b(T) > v_0^M$ which would also lead to incompatible final payoff and boundary data cannot happen by construction of the smoothed payoff.

# time steps	price
70	8.755
130	8.755
254	8.754
492	8.753

Table 1. Prices for DOP for different time steps using the Algorithm 1 with $\delta_{max} = 2$.

$w_{i_0} = u_{i_0}$. Then

$$\begin{aligned} w'_{i_0-1} &\leq w'_{i_0} + \delta(s_{i_0} - s_{i_0-1}) \\ &< w_{i_0-1} - \delta(s_{i_0} - s_{i_0-1}) + \delta(s_{i_0} - s_{i_0-1}) \end{aligned}$$

I.e. $i_0 - 1 \in Q$ which contradicts the minimality of i_0 . In a similar fashion one shows that the left sweep finds the smallest vector $v \geq w$ such that $v_i - v_{i-1} \leq \delta(s_i - s_{i-1})$. Moreover, the v_i also satisfy $v_i - v_{i-1} \geq -\delta(s_i - s_{i-1})$. This is clear if $v_{i-1} = v_i - \delta(s_i - s_{i-1})$. Otherwise $v_{i-1} = w_{i-1}$ and $v_i - v_{i-1} =$

$$v_i - w_{i-1} \geq w_i - w_{i-1} \geq -\delta(s_i - s_{i-1}).$$

This means that both sweeps together find the smallest vector $v \geq u$ such that $|v_i - v_{i-1}| \leq \delta(s_i - s_{i-1})$.

3 Results

For our numerical experiments we stick to the example of a DOP roughly 400 days to maturity where the current spot S is 76, the barrier is 62 (= 81%) and the strike 100.5 (= 132%) and a large dividend of 6% is paid in 15 days and a second one at T-15d. This was an extreme example in our development tests.

In our algorithm we use the well-known Rannacher scheme [6] to switch between implicit Euler and Crank-Nicholson time discretization, but algorithm BM is fully explicit with respect to the switch between Neumann and Dirichlet conditions. Thus it is worth looking at the convergence if the number of time steps increases. Table 1 shows prices for different time steps and a delta limit $\delta_{max} = 2$. The results in the table clearly show that algorithm BM does not deteriorate convergence properties.

In Table 2 we compare the prices of the classical management technique using shifted barriers with prices computed by our method. The barrier shift was determined so that the delta constraint is satisfied everywhere in the price surface for spots above the original barrier level and $t < T - 1day$. The

δ_{max}	classic (shifted bar. level)	algorithm BM
∞	6.51 (62.0)	6.51
8	8.28 (59.8)	6.79
4	12.16 (55.0)	7.19
2	14.47 (52.0)	8.76

Table 2. Price for DOP with different delta constraints for the simple barrier shift technique and Algorithm BM.

table shows a big price difference between the classical management technique and our method even for a relatively big delta limit $\delta_{max} = 8$ where the classical method gives a price of 8.29 compared to 6.79 using the generic approach.

In Figure 1 we show the complete price surfaces V with and without barrier management for the DOP. The δ_{max} parameter was 4. As we can see, the price difference is especially large in the region where the maximum delta occurs, i.e. short before maturity and close to the barrier. As we would have expected from the results in Table 2 the price difference is quite small at the beginning far away from the large discontinuity. Finally we investigate the performance hedging the DOP using the managed and unmanaged model. We apply simple delta hedging where the paths are generated with the local volatility dynamic which was used for pricing the DOP. This means that if we would hedge continuously we would end up with a final profit and loss with mean zero and zero variance. However, since we do the delta hedge only on a daily basis we would expect a non-zero variance. The barrier is watched continuously using an additional random number simulating the conditional minimum. To account for the gap risk the hedge is unwound at the end of the hedging interval if necessary.

It is important to emphasize that for the managed product each simulated P&L was reduced by the price offset to the unmanaged product, as we assume that we can not charge the price offset from the customer. To our opinion, although using the local volatility process to generate the paths on which we hedge is far away from being realistic, it is well suited to compare the managed and unmanaged strategies. The reason for this is that the hedge performance using the same model for backtesting as for pricing just shows the hedging errors coming from the discrete hedging what we are in our case mainly interested in.

Table 3 shows the 5% and 1% VaR and Ex-

	unmanaged	$\delta_{max} = 10$
5% VaR (ES)	-2.2 (-4.9)	-2.2 (-3.9)
1% VaR (ES)	-6.2 (-10.8)	-5.0 (-7.8)
0.1% VaR (ES)	-17.5 (-23.9)	-11.5 (-14.6)

Table 3. VaR and Expected Shortfall for the hedge backtest using no barrier management and the management with $\delta_{max} = 10$.

pected Shortfall (ES) for the hedge tests. The managed strategy reduces the 1% VaR and Expected Shortfall significantly although we are still using a quite large delta limit. To see an impact to the 5% VaR, the delta limit is still too big but nevertheless we see that the 5% Expected Shortfall is reduced by around 19%. Note that the initial cost to reduce the 1% expected shortfall by 3 was only around 0.26 which to our opinion is a very good trade off.

Figure 2 shows the histograms for the resulting P&L with and without the management technique. Looking at the P&L histogram for all paths, one can clearly see that the mean P&L for the managed product is shifted to the left compared to the unmanaged product. The reason for this is that the hedge costs for the managed product are higher which is reflected by the higher price, but we get just the lower price of the unmanaged product at the beginning. This is the general disadvantage of using management techniques: If the barrier is not hit at the beginning, one has lost money hedging a more expensive product. On the other hand, the down side risk is massively reduced. In addition to the risk measures presented in Table 3 this can be clearly seen in Figure 2 where also the histograms of the P&L for all paths hitting the barrier are plotted.

4 Barrier management for Monte-Carlo products

A barrier management technique similar to the finite difference case can be easily implemented in the Monte-Carlo payoff smoothing technique developed in [1] for auto-callable instruments. These instruments have discrete monitoring dates T_1, \dots, T_m called redemption dates.

For each redemption date a trigger function U is evaluated at the simulated spots $S^\omega(T_k) := (S_1^\omega(T_k), \dots, S_d^\omega(T_k))^t$. Here, d is the number of underlyings and ω is the path index. If the instrument was not called before, it is called at T_k if

$U(S^\omega(T_k)) \geq$ redemption level R . In this case the owner receives a notional N . If the instrument was not called until T_m , the owner receives a final payoff $F_m(U(S^\omega(T_m)))$. For equity derivatives F_m is often a piece-wise linear function with jumps.

The digital features make it hard to compute stable greeks using finite differences with prices from shifted market data. For that a payoff-smoothing like technique was developed in [1] which extends prior work of [2], [3], [5]. In the following we give a brief description of this method. For reasons of simplicity we drop discounting cash flows.

First each path is simulated on a time grid which contains the redemption dates with whatever model (local or stochastic volatility) and whatever scheme (Euler, log Euler, Milstein) one wants.

The path value is computed from back to front similar to a PDE solver. For the final redemption date T_m we pick a simulated time point $T'_m < T_m$. Usually the simulated time point right before T_m . We assume that the spots move locally like a shifted Brownian motion starting in $S^\omega(T'_m)$, i.e. there is a $d \times d$ matrix L such that

$$S(T_m) = S^\omega(T'_m) + L \cdot W, \quad W \sim N(0, id).$$

Now, one introduces a coordinate transform by means of an orthogonal matrix Q . The last column h must be such that the function $u(S, x) := U(S + Lhx)$ is strictly monotone w.r.t. x in a sufficiently large neighbourhood of $(S^\omega(T'_m), 0)$. This sounds complicated, but is easily verified for e.g. $h = L^{-1}(1, \dots, 1)$ and common trigger functions like worst-of or best-of. To compute the path value one draws $d - 1$ independent Gaussian random numbers X_1, \dots, X_{d-1} and puts them in a d -dimensional vector $X = (X_1, \dots, X_{d-1}, 0)$. Then, the path value is initialized by

$$V_m^\omega := E_{X_d}[F_m(U(S^\omega(T'_m) + LQX + LQe_d X_d))] \quad (4)$$

where e_d is the vector with zeros at coordinates $1, \dots, d - 1$ and equal to one at the d -th coordinate. This expected value is computed semi-analytically or numerically. The method is a mixture of Monte-Carlo sampling – because of X – and local semi-analytical treatment. Details how to choose e.g. h , Q and an analysis of the bias can be found in [1].

For other redemption dates $T_k, k < m$ the path value is updated in a similar fashion. We already know V_{k+1}^ω . Again, we pick $T'_k < T_k$ and again assume that there is a matrix L such that

$$S(T_k) = S^\omega(T'_k) + L \cdot W, \quad W \sim N(0, id)$$

and also compute h, R and X as before. The updated path value then is

$$V_k^\omega := E_{X_d}[F_k(U(S^\omega(T'_k) + LQX + LQe_n X_d))] \quad (5)$$

where

$$F_k(u) := \begin{cases} N & \text{if } u \geq R \\ V_{k+1}^\omega & \text{else} \end{cases}.$$

The final value of the path is V_1^ω .

With these prerequisites it is quite easy to introduce the barrier-management technique: one simply replaces F_k in (4) and (5) by smoothed variants \tilde{F}_k which fulfill the constraints

$$\begin{aligned} |\tilde{F}'_k(u)| &\leq \delta_{max} \\ \tilde{F}(u) &\geq F(u). \end{aligned}$$

By that we have that the delta of the managed instrument can not be larger than $\delta_{max} |\nabla U|_\infty$.

As an example we consider an auto-callable on three underlyings with four redemption dates, the first being one year ahead, the second two years and so on.

The trigger function is a worst-of three single stocks defined relative to the spots at $t = 0$ as we assume that the forward start period is just over.

$$U(S_1, \dots, S_3) = \min(S_1/S_1(0), \dots, S_3/S_3(0)).$$

The redemption level is always 1. The redemption notional is 1 and the final payoff is

$$F_4(u) = \begin{cases} 1 & \text{if } u \geq 0.6 \\ u & \text{else} \end{cases}$$

This test case resembles the situation when the instrument has just been issued and is traded for the first time. Table 4 shows prices and deltas for the plain Monte-Carlo scheme, Monte-Carlo payoff-smoothing and Monte-Carlo payoff-smoothing with different levels of δ_{max} .

The second test case considers a similar instrument that has been issued 5 years ago and the final redemption date is just 4 hours ahead. The instrument has not been called before and two spots have moved down by 20% and one by 40% such that the trigger function value is right in front of the jump at 60% in the final payoff. The results presented in Table 4 show the impact of the management technique for different maximum deltas on the price and deltas of the product.

	price	$\delta_i \cdot S_i(0), i = 1, \dots, 3$		
MC	0.7389	0.2375	0.2728	0.2855
PS	0.7392	0.2376	0.2710	0.2834
$\delta = 20$	0.7470	0.2383	0.2731	0.2848
$\delta = 10$	0.7557	0.2375	0.2715	0.2847
$\delta = 5$	0.7756	0.2348	0.2683	0.2760

Table 4. Price and deltas for a worst-off auto-callable for the unmanaged Monte Carlo method (MC), with payoff smoothing (PS) and the managed methods with different maximum deltas.

	price	$\delta_i \cdot S_i(0), i = 1, \dots, 3$		
MC	0.7984	0	0	19.484
PS	0.7986	0	0	19.481
$\delta = 20$	0.9766	0	0	6.000
$\delta = 10$	0.9883	0	0	3.000
$\delta = 5$	0.9941	0	0	1.500

Table 5. Price and deltas for a worst-off auto-callable very close to final redemption at trigger level for the unmanaged Monte Carlo method (MC), with payoff smoothing (PS) and the managed methods with different maximum deltas.

5 Conclusion

We have presented a generic algorithm which can easily be implemented in a PDE pricing framework to take the digital risk of a product into account for pricing and hedging. Here, one can simply control the gap risk by setting a delta limit which may be determined by the liquidity of the traded underlying. Since we are modifying the solution at certain time slices within the PDE solution method, one may raise the question how this process is influenced by the time step size. Here, numerical experiments indicate that it is robust w.r.t. to time step size. It turned out that the method presented in this paper produces much lower prices than the classical approach by simply shifting the barrier. This is an important feature since this price difference directly reflects the cost of hedging the managed product and can be interpreted as the premium paid to get insurance against the gap risk. In this sense one has to pay a lower premium using the generic method compared to the simple shifting methodology.

To show the impact using this management technique we presented some hedging backtests for a DOP showing a significant reduction of VaR and Expected Shortfall compared to the unmanaged strategy.

Finally we discussed how this technique may be carried over to Monte Carlo pricing for auto-callable products and showed promising results for a two dimensional structure with four redemption dates. However, how to deal with barrier management in the MC setting in a more general way such as for continuous barriers is still open and remains future work.

References

- [1] *Monte-Carlo Payoff-Smoothing for Pricing Auto-callable Instruments* F.Koster and A.Rehmet, Available at SSRN: <http://ssrn.com/abstract=2602905>
- [2] *Conditional analytic Monte-Carlo pricing scheme of auto-callable products* C.P. Fries and M.S. Joshi, Available at SSRN: <http://ssrn.com/abstract=1125725>
- [3] *A Monte-Carlo pricing algorithm for auto-callables that allows for stable differentiation* T. Alm, B. Harrach, D. Harrach and M. Keller, J. Comput. Finance, Vol. 17.2013, 1, p. 43-70
- [4] P. Glasserman *Monte-Carlo methods in financial engineering*; vol. 53 of Stochastic Modelling and Applied Probability, Springer 2003.
- [5] P. Glasserman and J. Staum *Conditioning on one-step survival for barrier option simulations*; Operations Research 49(6), 923937, 2001.
- [6] *Finite Element solution of diffusion problems with irregular data* R. Rannacher, Numerische Mathematik, vol. 43, p. 309-327




Figure 1. DOP: price surfaces without barrier management (left), $\delta_{max} = 4$ (middle) and their difference (right).




Figure 2. P&L distribution for the DOP for unmanaged security (left) and managed (right). The bottom row is a close up of the top row with focus on rare events.