OPTION PRICING AND HEDGING WITH TEMPORAL CORRELATIONS

Lorenzo Cornalba[1], Jean–Philippe Bouchaud[2,3] and Marc Potters[3]

[1] Laboratoire de Physique Théorique de l'École Normale Supérieure
24 rue Lhomond, 75231 Paris Cedex 5, France
[2] Service de Physique de l'État Condensé, Centre d'études de Saclay
Orme des Merisiers, 91191 Gif–sur–Yvette Cedex, France
[3] Science & Finance
The Research division of Capital Fund Management
109–111 rue Victor-Hugo, 92532 Levallois Cedex, France

April 9, 2001

Abstract

We consider the problem of option pricing and hedging when stock returns are correlated in time. Within a quadratic–risk minimisation scheme, we obtain a general formula, valid for weakly correlated non–Gaussian processes. We show that for Gaussian price increments, the correlations are irrelevant, and the Black–Scholes formula holds with the volatility of the price increments on the scale of the re–hedging. For non–Gaussian processes, further non trivial corrections to the 'smile' are brought about by the correlations, even when the hedge is the Black–Scholes Δ–hedge. We introduce a compact notation which eases the computations and could be of use to deal with more complicated models.

1 Introduction

The assumptions underlying the Black–Scholes model [1, 2] – that the dynamics of financial assets can be modelled by a continuous time Gaussian process – are very far from reality. It is well known, in particular, that the empirical distribution of returns exhibits 'fat–tails'. The presence of these extreme events not only induces a 'smile' in the implied volatility of the options, but also destroys the essence of the Black–Scholes pricing procedure: the possibility of constructing a riskless hedge which replicates perfectly the option pay–off [3, 4, 5]. A more subtle further effect, completely discarded by the standard Black–Scholes model, is the presence of small, but significant, temporal *anticorrelations* in the series of stock returns. Correspondingly, the variance of the N–day distribution of returns is smaller than N times the variance of the daily distribution (see *e.g.* [6]). On other financial assets, such as stock indices, one observes the opposite effect of weak positive correlations. The order of magnitude of these correlations or anticorrelations is 10% from one day to the next. These correlations lead, in principle, to the possibility of statistical arbitrage. However, one should also take into account transaction costs which significantly reduce the possibility of using these correlation in practice [5]. In other words, the presence of transaction costs allow the existence of non zero correlations, which can be significant and therefore must be dealt with consistently in the context of option pricing.

What is the influence of these correlations on the price of an option and on the corresponding hedging strategy? The aim of this paper is to provide a general answer to this question within a quadratic risk minimization sheme, in the case where these correlations are small [1] (which is, in most cases of interest, amply justified), with only very few assumptions on the actual distribution of returns. We calculate the first order corrections both to the price and optimal hedge. In the Gaussian case, we find that the effect of correlations can be compensated by a change in the hedging strategy and therefore options should be priced using the standard uncorrelated Black–Scholes model [2]. The correct volatility to be used is the one measured on the time scale of the rehedging, and not the one corresponding to the terminal distribution – *i.e.* corresponding to the maturity of the option. In other words, the 'risk–neutral' measure that one should use to price derivatives is simply an uncorrelated Gaussian measure with the same elementary variance. These conclusions are not valid if one considers a process for which the underlying distribution is non–Gaussian. One then finds contributions to the option prices which depend in general on the strength of the correlations, and on the higher moments of the distribution (in particular on the kurtosis). The corrections to the option price also depend on the past history of the price–increments of the underlying. The detailed analysis of a specific example shows, in particular, that anticorrelations change the effective kurtosis to be used when computing 'smile' corrections to the options prices. In general, the correct kurtosis to be used is neither the local nor the global one, but some effective kurtosis which depends on the strength of the correlations.

[1] The general correlated Gaussian case was considered in [3], and leads to rather complex formulae for which no simple interpretation was given.

[2] This can be easily seen within a Black–Scholes framework: see Appendix.

2 General setting

Our starting point is the global wealth balance approach discussed in [3, 5]. We set, for simplicity, the interest rate to zero, and consider a discrete time model with time interval τ, within which no rehedging takes place. The global wealth balance at time $T = N\tau$ for the writer of the option reads

$$W = \mathcal{C} + H - \Theta,$$

where $\mathcal{C}$ is the option premium, Θ is the pay–off of the option, and H is the result of the hedging strategy

$$H = \sum_{n=0}^{N-1} \phi_n \Delta_{n+1}. \tag{1}$$

In equation (1), $\Delta_{n+1} = x_{n+1} - x_n$ denotes the change of the price x of the underlying asset between time n and $n+1$, and ϕ_n is the amount of underlying in the portfolio at time n [3].

In incomplete markets, the price of an option is not unique since a risk–premium should be added. However, a convenient framework is the risk minimisation scheme proposed in [3, 4, 5], that fixes both the price and strategy $\mathcal{C}$, ϕ_n so as to minimize the total risk $\mathcal{R}$ associated to option writing, which we simply define as

$$\mathcal{R}^2 = \langle W^2 \rangle ,$$

where $\langle ... \rangle$ is an average over the objective probability at time 0. Minimizing with respect to $\mathcal{C}$ fixes the price of the call option to be

$$\mathcal{C} = \langle \Theta - H \rangle . \tag{2}$$

If, on the other hand, we minimize $\mathcal{R}^2$ with respect to ϕ_n, we obtain an equation implicitly determining the optimal strategy $\phi_n^\star$, equation that reads, for $0 \leq n < N$

$$\langle \Delta_{n+1} (\Theta - H) \rangle_n = \mathcal{C} \langle \Delta_{n+1} \rangle_n \tag{3}$$

The notation $\langle ... \rangle_n$ in the above equation means that the average (again over the objective probability) is performed at time n – *i.e.* with all the information available at this instant of time. This notation turns out to be very powerful, and could be used to handle very general situations. We now analyze the above general and compact equations (2) and (3) in several simple but important limiting cases.

[3] In the following, we assume for simplicity that the interest rate is zero. It is easy to generalize the formulae obtained here when this rate is non-zero: see [5] for details.

3 The case of zero conditional drift

Let us define the conditional drift μ_n and the conditional variance D_n at time n as

$$\begin{aligned} \mu_n &= \langle \Delta_{n+1} \rangle_n \\ D_n &= \left\langle (\Delta_{n+1} - \mu_n)^2 \right\rangle_n . \end{aligned}$$

Note that these quantities are the expected drift and variance for the next time interval, and depend in general on the past (realized) history of price changes.

We start by assuming that the drift μ_n is identically zero (this is to say, the price process is a 'martingale'). Then, since ϕ_n is known at time n, one has that

$$\langle \phi_n \Delta_{n+1} \rangle = \langle \phi_n \langle \Delta_{n+1} \rangle_n \rangle = 0,$$

and therefore that $\langle H \rangle = 0$. This implies that the option price is nothing but the expectation value of the pay–off

$$\mathcal{C} = \langle \Theta \rangle .$$

To determine the optimal hedge, we note that, in general, one has

$$\begin{aligned} \langle \Delta_{n+1} H \rangle_n &= \sum_{m=0}^{N-1} \langle \phi_m \Delta_{m+1} \Delta_{n+1} \rangle_n \\ &= \phi_n \left\langle \Delta_{n+1}^2 \right\rangle_n + \sum_{m=0}^{n-1} \phi_m \Delta_{m+1} \mu_n + \sum_{m=n+1}^{N-1} \langle \phi_m \mu_m \Delta_{n+1} \rangle_n . \end{aligned} \tag{4}$$

In the case of vanishing drift $\mu_n = 0$ we therefore have that

$$\langle \Delta_{n+1} H \rangle_n = \phi_n D_n$$

so that the optimal hedging strategy can be computed, from equation (3), as [3, 5]

$$\phi_n^\star = \frac{1}{D_n} \langle \Delta_{n+1} \Theta \rangle_n .$$

4 Perturbative expansion for small drift

Let us now turn to the case where the conditional bias μ_n is non–zero but small. Using equation (4), it is easy to see that the basic equations (2) and (3), which determine the optimal price and the hedging strategy, can be solved perturbatively in μ_n. For our purposes, we will only need the first correction in the conditional drift, which we now derive. First note that equation (2) can be rewritten as

$$\mathcal{C} = \langle \Theta \rangle - \sum_{n=0}^{N-1} \langle \phi_n^* \mu_n \rangle .$$

In particular, to first order in μ_n, one finds that

$$\mathcal{C} = \langle \Theta \rangle - \sum_{n=0}^{N-1} \left\langle \frac{\Delta_{n+1}\mu_n}{D_n} \Theta \right\rangle . \tag{5}$$

One can also compute the first correction to the hedging strategy, by combining equation (4) with the zero–th order results obtained in the last section. One obtains then

$$\phi_0^\star = \frac{1}{D_0} \langle \Delta_1 \Theta \rangle - \frac{1}{D_0} \sum_{n=1}^{N-1} \left\langle \frac{\Delta_{n+1}\mu_n}{D_n} \Delta_1 \Theta \right\rangle - \frac{1}{D_0} \mu_0 \langle \Theta \rangle . \tag{6}$$

In the above equation, we have only shown the optimal hedge for the first time period $\phi_0^\star$. The explicit expression for the optimal hedge for the subsequent time periods is also of a similar form, but it is quite intricate and not very illuminating. More importantly, it is not necessary to determine the correct hedging strategy. This important point will be discussed, in a slightly more general context, in section 7.

5 A general model for weakly correlated processes

Let us introduce a rather general model for correlated price increments. We first define an auxiliary set of *uncorrelated i.i.d. random variables* $\{\eta_n\}$, distributed according to an arbitrary probability distribution[4] $P(\eta)$. The joint distribution is then given by

$$P(\{\eta_n\}) = \prod_n P(\eta_n).$$

We will assume that $P(\eta)$ is symmetrical[5] – *i.e.* that

$$\int \eta^{2q+1} P(\eta) d\eta = 0, \tag{7}$$

[4]The following results can be generalized to the case where the distribution $P(\eta)$ explicitly depends on n.

[5]In the sequel, we will actually only need that equation (7) holds for $q = 0,1$.

and we will denote the second moment of P by

$$D = \int \eta^2 P(\eta) d\eta.$$

We now construct the set of correlated price increments $\{\Delta_n\}$ by writing

$$\Delta_n = \sum_m M_{nm} \eta_m + \mu, \tag{8}$$

with

$$M_{nm} = \delta_{n,m} + \frac{1}{2D} c_{n-m}.$$

The coefficients c_n are assumed to satisfy

$$c_0 = 0 \qquad\qquad c_{-k} = c_k.$$

In the sequel, we assume that both the c_k's and μ are small, and will work only to first order in both μ and c.

Let us show the significance of equation (8), by first noting that, since $\langle \eta_n \rangle = 0$, we have that

$$\langle \Delta_n \rangle = \mu.$$

Therefore μ is nothing but the average *unconditional* drift of the price process. Moreover, to first order in c, the following holds

$$\langle (\Delta_n - \mu)(\Delta_m - \mu) \rangle = \sum_{i,j} M_{ni} M_{mj} \langle \eta_i \eta_j \rangle = D\delta_{n,m} + c_{n-m}.$$

The coefficients D and c denote then, respectively, the *unconditional* variance and correlation of the price increments. Finally, one has that, independently of c_{n-m}, and for all p,

$$\langle (\Delta_n - \mu)^p \rangle = \langle \eta^p \rangle .$$

We note that the variance of the price differences over a time scale n, given by

$$\sigma^2(n) = \left\langle \left(\sum_{k=m+1}^{m+n} (\Delta_k - \mu) \right)^2 \right\rangle , \qquad (n \geq 1)$$

is related to the correlation coefficients c_k as follows

$$\begin{aligned} \sigma^2(n) &= nD + \sum_{k=1}^{n-1} 2(n-k)c_k \\ c_n &= \frac{1}{2}\left(\sigma^2(n+1) - 2\sigma^2(n) + \sigma^2(n-1)\right). \end{aligned} \tag{9}$$

In the case where $c_n = 0$, one finds $\sigma^2(n) = nD$, whereas, in the case of very short range anticorrelations ($c_n = -c\delta_{n,1}$), one has $\sigma^2(n) = n(D - 2c) + 2c$. In this case, the global volatility $\sigma^2(n)/n \to D - 2c$ is lower than the short–range volatility $\sigma^2(1) = D$.

The probability distribution of the price increments can be written, to first order in μ, c, as

$$\begin{aligned} &\prod_n P\left(\Delta_n - \mu - \frac{1}{2D}\sum_m c_{n-m}\Delta_m\right) \\ = &\left[1 - \mu\sum_n \frac{\partial \ln P}{\partial \Delta_n} - \frac{1}{2D}\sum_{n,m}\frac{\partial \ln P}{\partial \Delta_n}c_{n-m}\Delta_m\right]\prod_n P(\Delta_n). \end{aligned} \tag{10}$$

It is now easy to write the marginal distribution at time $n = 0$, given all previous returns Δ_n with $n \leq 0$

$$\left[1 - \mu\sum_{n>0}\frac{\partial \ln P}{\partial \Delta_n} - \frac{1}{2D}\sum_{\max(n,m)>0}\frac{\partial \ln P}{\partial \Delta_n}c_{n-m}\Delta_m\right]\prod_{n>0} P(\Delta_n). \tag{11}$$

The equation above follows immediately from the general expression for the probability distribution, up to a possible normalization factor. To show that (11) is actually correctly normalized, we have only to note the obvious fact that

$$\int d\Delta\, P(\Delta)\frac{\partial \ln P}{\partial \Delta} = 0.$$

Moreover, using that

$$\begin{aligned} \int d\Delta\, P(\Delta)\Delta\frac{\partial \ln P}{\partial \Delta} &= -1 \\ \int d\Delta\, P(\Delta)\Delta^2\frac{\partial \ln P}{\partial \Delta} &= 0 \end{aligned}$$

it is then easy to show that the *conditional* drift and variance μ_n and D_n are given, to first order in μ and c, by

$$\begin{aligned} \mu_{n-1} &= \mu + \sum_{m<n} c_{n-m}\left(\frac{1}{2D}\Delta_m - \frac{1}{2}\frac{\partial \ln P}{\partial \Delta_m}\right) \\ D_n &= D. \end{aligned}$$

6 Derivative pricing with small correlations

6.1 A general formula

Let us now apply the general perturbative formula (5) for the price of the contract to the specific model of correlations described in the previous section . One finds

$$\begin{aligned}\mathcal{C} &= \langle \Theta \rangle - \frac{1}{D} \sum_{n=1}^{N} \langle \mu_{n-1} \Delta_n \Theta \rangle \\ &= \langle \Theta \rangle - \frac{\mu}{D} \sum_{n=1}^{N} \langle \Delta_n \Theta \rangle - \frac{1}{2D} \sum_{n=1}^{N} \sum_{m<n} c_{n-m} \left\langle \left(\frac{1}{D} \Delta_m - \frac{\partial \ln P}{\partial \Delta_m} \right) \Delta_n \Theta \right\rangle .\end{aligned}$$

Now, the above averaging $\langle \cdots \rangle$ is performed over the 'correlated' historical price increment distribution, equation (11). It is convenient to re–express all of the following formulæ in terms of an auxiliary probability distribution $\prod_{n>0} P\left(\Delta_n\right)$, which describes an unbiased and uncorrelated process, and is obtained by setting $c_k = \mu = 0$. We will denote the corresponding averages by $\langle ... \rangle^0$, where the superscript indicates that the conditional drift is now set to zero (in the case of a risk–free Gaussian model, this auxiliary probability distribution is called the equivalent martingale measure in the financial mathematics literature). Equation (11) can then be reinterpreted as relating the two averaging procedures $\langle \cdots \rangle$ and $\langle \cdots \rangle^0$ as follows

$$\langle \cdots \rangle = \langle \cdots \rangle^0 - \mu \sum_{n>0} \left\langle \cdots \frac{\partial \ln P}{\partial \Delta_n} \right\rangle^0 - \frac{1}{2D} \sum_{\max(n,m)>0} \left\langle \cdots \frac{\partial \ln P}{\partial \Delta_n} c_{n-m} \Delta_m \right\rangle^0 .$$

After a some algebra we then find that

$$\mathcal{C} = \langle \Theta \rangle^0 + \mu \sum_{n=1}^{N} \langle F\left(\Delta_n\right) \Theta \rangle^0 + \frac{1}{2D} \sum_{n=1}^{N} \sum_{m<n} c_{n-m} \langle F(\Delta_n) \Delta_m \Theta \rangle^0 , \qquad (12)$$

where

$$F(\Delta) = -\frac{\partial \ln P}{\partial \Delta} - \frac{1}{D} \Delta .$$

The above equation is the central result of this paper, that we now comment in various limits.

Let us first consider the purely Gaussian case, where $\ln P(\Delta) = -\Delta^2/2D - 1/2 \ln(2\pi D)$. In this situation, it is easy to check that $F(\Delta) = 0$. Therefore, all corrections brought about by the drift and correlations strictly vanish, and the price can be calculated by assuming that the process is correlation–free. As shown in the Appendix, this is indeed expected in a continuous time framework. In other words, the price of the option is given by the Black–Scholes price, with a volatility given by the small scale volatility D (that corresponds to the hedging time scale), and *not* with the volatility corresponding to the terminal distribution, $\sigma^2(N)/N$. As we have already shown, when price increments are anticorrelated, D is larger than $\sigma^2(N)/N$, and the price of the contract is higher than the objective average pay–off. This is due to the fact that the hedging strategy is on average losing money because of the anticorrelations: a rise of the price of the underlying leads to an increase of the hedge, which is followed (on average) by a drop of the price.

For an uncorrelated non–Gaussian process, the effect of a non–zero drift μ does not vanish, a situation already discussed in details in [5]. The same is true with correlations: the correction no longer disappears and one cannot price the option by considering a process with the same statistics of price increments but no correlations. Note that the correction term involves all Δ_m with $m < 0$. This means, as expected a priori, that, in the presence of correlations, the price of the option does not only depend on the current price of the underlying, but also on all past price increments.

6.2 The case of short range correlations

Let us illustrate the general formula in the case of very short range anticorrelations $c_n = -c\delta_{n,1}$. We will furthermore consider the simple case of path–*independent* pay–offs Θ, which are functions only on the terminal price x_N of the underlying.

The calculation of the price will therefore involve the following quantities

$$I_1 = \langle F\left(\Delta_n\right)|x_N\rangle^0 \qquad\qquad I_2 = \langle \Delta_n F\left(\Delta_m\right)|x_N\rangle^0,$$

where the notation $\langle ...|x_N\rangle^0$ means that we are conditioning on a given value of the terminal price x_N. The computation of the two quantities I_1, I_2 is easily done in momentum space, with the aid the following formulae

$$\begin{aligned}
\int d\Delta\ \Delta P(\Delta) e^{iz\Delta} &= iDz - \frac{i}{6}D^2\left(3+\kappa\right)z^3 + o\left(z^5\right)\\
\int d\Delta\ F\left(\Delta\right)P(\Delta)\,e^{iz\Delta} &= \frac{i}{6}\kappa D\,z^3 + o\left(z^5\right),
\end{aligned}$$

where κ is the kurtosis of the distribution of elementary increments. The result reads, for large N,

$$I_1 = -\frac{\kappa}{N\sqrt{DN}}p_1 \qquad\qquad I_2 = -\frac{4\kappa}{N^2}p_2,$$

where p_1 and p_2 are, respectively, the skewness and kurtosis polynomials

$$p_1 = \frac{1}{6}\left(\widetilde{x}^3 - 3\widetilde{x}\right) \qquad\qquad p_2 = \frac{1}{24}\left(\widetilde{x}^4 - 6\widetilde{x}^2 + 3\right)$$

and where we have consider the natural scaling $\widetilde{x}^2 = \left(x_N - x_0\right)^2/DN$: $\widetilde{x}$ is the moneyness of the option counted in standard deviations. The general equation (12) then takes the form $\mathcal{C} = \langle\widetilde{\Theta}\rangle^0$, with a modified pay–off function

$$\widetilde{\Theta}(x_N) = \Theta(x_N)\left[1 - \frac{\kappa}{\sqrt{DN}}\left(\mu - \frac{c\Delta_0}{2DN}\right)p_1 + \frac{2c\kappa}{DN}p_2\right],$$

where the explicit dependence on the last past increment Δ_0 appears. Note that the correction terms vanish for at the money options ($x_N = x_0$). Let us concentrate on the last term in the above equation. It represents an *increase* in the global kurtosis κ/N of the 'risk neutral' distribution that must be used to price the option:

$$\frac{\kappa}{N} \to \frac{\kappa}{N}\left(1 + \frac{2c}{D}\right). \tag{13}$$

This is means that the 'implied' kurtosis is increased, in percent terms, as much as the global variance is decreased from the local one by correlation effects. Therefore the volatility 'smile', which is proportional to the kurtosis, is enhanced by the presence of anti–correlations, and reduced by the presence of correlations. Note that the above model neglects all volatility fluctuations, which lead to a terminal kurtosis decaying much more slowly than κ/N (see [5]). The above calculations could in principle be extended to deal with this effect, which is crucial for financial applications.

6.3 An alternative model of correlations

Let us comment more on the above result. To this end, let us briefly discuss an other possible model of correlated price increments. One could consider, instead of the probability distribution (10), the following distribution for price increments (we set for simplicity $\mu = 0$)

$$\left[1 + \frac{1}{2D^2}\sum_{n,m} c_{n-m}\Delta_n\Delta_m\right]\prod_n P(\Delta_n). \tag{14}$$

The above measure induces statistics similar to (10), and in particular one still has that $\langle\Delta_n\Delta_m\rangle = D\delta_{n,m} + c_{n-m}$, and that $\langle\Delta_n^p\rangle = \langle\eta^p\rangle$. On the other hand, it is not hard to show, following an analysis similar to the one in section 3, that, to first order in the correlations, the price of an option is *unchanged,* regardless of the basic day–to–day probability distribution $P(\eta)$. How can one decide which model more accurately describes a specific time–series of price increments? To answer this question, we first note that (10) and (14) coincide when $P(\eta)$ is Gaussian. It must then be that the deviations involve the higher moments of P. Let us denote with λ_i the i–th cumulant of P (in particular $\lambda_2 = D$ and $\lambda_4 = D^2\kappa$, with κ the kurtosis), and with $\lambda_i(n)$ the i–th cumulant of the return $\Delta_1 + \cdots + \Delta_n$ of the stock in n time periods. If the price increments where independent, then $\lambda_i(n) = n\lambda_i$. In the presence of correlations, this equation is not valid any more, but, if the correlations are short–range, we can define global cumulants as

$$\widetilde{\lambda}_i = \frac{1}{n}\lambda_i\,(n)\,. \qquad\qquad (n \to \infty)$$

To be definite, let us consider the case already discussed of short–range anticorrelations $c_1 = -c$. We have already shown that $\lambda_2\,(n) = n\,(D - 2c) + 2c$ and therefore that

$$\widetilde{\lambda}_2 = D\left(1 - \frac{2c}{D}\right).$$

It is then not hard to show that, if one considers the model (10), then the global forth moment is given by

$$\widetilde{\lambda}_4 = D^2\kappa\left(1 - \frac{4c}{D}\right).$$

On the other hand, if one uses (14) one finds that

$$\widetilde{\lambda}_4 = D^2\kappa\left(1 - \frac{8c}{D}\right).$$

In both cases, the global historical kurtosis is reduced by anti–correlation effects, but the amount varies in the two models. Recalling that $2c/D \simeq 0.10$, we see that the two models differ by 10% in the measurement of the global kurtosis $\widetilde{\lambda}_4/D^2\kappa$.

6.4 The case of Delta-hedging

Le us conclude this section by discussing other possible hedging schemes. Instead of choosing the optimal hedge that minimizes the variance, one can follow the (sub–optimal) Black–Scholes Δ–hedge. As emphasized in [5], this leads in general to a negligible risk increase, and has the advantage of removing the effect of the drift μ on the price, and therefore reduces the risk associated to change of long time trends, encoded in μ. One might therefore expect that Δ–hedging also allows one to get rid of the correlations. However, the following calculation shows that this is not the case.

Up to first order in c, it is sufficient to consider that the Δ–hedge is given by

$$\phi_{BS,n} = \frac{\partial}{\partial x_n}\langle\Theta\rangle_n^0 = -\langle\frac{\partial \ln P}{\partial\Delta}(\Delta_{n+1})\Theta\rangle_n^0.$$

The price of the contract is then given by

$$
\begin{aligned}
\mathcal{C} &= \langle\Theta\rangle - \sum_{n=0}^{N-1} \langle \mu_n \phi_{BS,n} \rangle \\
&= \langle\Theta\rangle + \sum_{n=0}^{N-1} \left\langle \mu_n \frac{\partial \ln P}{\partial \Delta_{n+1}} \Theta \right\rangle \\
&= \langle\Theta\rangle^0 + \frac{1}{2} \sum_{n=1}^{N} \sum_{m<n} c_{m-n} \langle F(\Delta_n) \frac{\partial \ln P}{\partial \Delta_m} \Theta \rangle^0 .
\end{aligned}
$$

Note that, in this case, the drift μ indeed disappears to first order even if $F(\Delta)$ is not zero; the correlations, on the other hand, are still present within a Δ–hedging scheme.

7 Other optimization schemes

Let us now consider a slightly different problem, where the price of the option is fixed by the market to a certain value $\mathcal{C}_M$. The option trader knows that price increments have a non–zero conditional drift and wishes to buy and hedge the option in order to optimize a certain risk adjusted return. The wealth balance associated to selling the option now reads

$$W = \mathcal{C}_M + H - \Theta,$$

so that the expected return is equal to

$$M = \mathcal{C}_M + \langle H - \Theta \rangle = \mathcal{C}_M - \mathcal{C}$$

and the expected risk is

$$\mathcal{R}^2 = \left\langle W^2 \right\rangle - M^2 .$$

We will assume that the trader wants to maximize a certain function G of M and $\mathcal{R}$, for example the Sharpe ratio $G_1 = M/\mathcal{R}$, or a certain risk corrected return $G_2 = \lambda M - \mathcal{R}$.

Consider a small variation $\delta\phi_n$ of the strategy ϕ_n. The corresponding changes in M and $\mathcal{R}$ read

$$\frac{\delta M}{\delta \phi_n} = \mu_n \qquad \mathcal{R} \frac{\delta \mathcal{R}}{\delta \phi_n} = \langle \Delta_{n+1} (H - \Theta) \rangle_n + \mathcal{C}\, \mu_n .$$

We may now extremize G with respect to the hedging strategy and obtain the following equation

$$\frac{\delta M}{\delta \phi_n} \frac{\partial G}{\partial M} + \frac{\delta \mathcal{R}}{\delta \phi_n} \frac{\partial G}{\partial \mathcal{R}} = 0,$$

or

$$\langle \Delta_{n+1}(\Theta - H)\rangle_n = \mu_n \mathcal{C} + \mu_n \mathcal{G} \qquad \mathcal{G} = \mathcal{R}\frac{\partial_M G}{\partial_{\mathcal{R}} G}. \tag{15}$$

In the two examples considered before, one has that

$$\mathcal{G}_1 = -\frac{\mathcal{R}^2}{M} \qquad \mathcal{G}_2 = -\lambda \mathcal{R}.$$

Equation (15) is a generalization of equation (3), and is again well suited for a perturbation in μ. We will work as always to first order in μ, and we will call $\widetilde{\phi}$ the solution to (15) with the function $\mathcal{G}$ set to zero – that is, the solution to the problem solved in section 4). It is then easy to see, using equation (4), that the optimal strategy is given by

$$\phi_n^* = \widetilde{\phi}_n - \frac{\mu_n}{D_n}\, \mathcal{G}|_{\widetilde{\phi}}\,. \tag{16}$$

We conclude this section by showing that the more general formalism developed above actually has some interesting consequences for the original problem of option pricing and hedging which was considered in the rest of the paper. Let us in particular consider, as we have done in the previous sections, a specific contract sold at time 0 at the optimal price $\mathcal{C}_0$ and with optimized hedging strategy $\phi^\star_{n,0}$ (we will explicitly insert the time–index for the date of *sale* of the contract in the last part of this section). More generally, we may sell the same contract at some other time P at a price $\mathcal{C}_P$ with a hedging strategy $\phi^\star_{n,P}$ (with $n \geq P$). It is a natural question to ask what is the relation of the various hedging strategies $\phi^\star_{n,P}$ for *fixed* time n, as we vary the sale time of the contract P. In particular, one may ask if the optimal hedging strategy $\phi^\star_{n,P}$ at time n should be equal to the optimal hedge $\phi^\star_{n,n}$ which one is to adopt if one is selling the contract at time n. This question can be easily answered in the framework of the present section, by noting that, at time n, the option traders wealth balance is given by the original price of the contract $\mathcal{C}_P$ plus the returns of the hedging strategy up to time n

$$\mathcal{W}_n = \mathcal{C}_P + \phi^\star_{P,P}\Delta_{P+1} + \cdots + \phi^\star_{n-1,P}\Delta_n.$$

From the point of view of the trader, the above quantity acts as an effective market price $\mathcal{C}_M$, and the optimization problem after time n is of the type described at the beginning of this section. Recalling that the trader is trying to minimize $\langle W^2\rangle$, we deduce that the function G is given by $M^2 + \mathcal{R}^2$, and therefore that $\mathcal{G} = M = \mathcal{C}_M - \mathcal{C} = \mathcal{W}_n - \mathcal{C}_n$. We can then use equation (16) to deduce that

$$\phi^\star_{n,P} = \phi^\star_{n,n} - \frac{\mu_n}{D_n}\left(\mathcal{W}_n - \mathcal{C}_n\right).$$

The above equation has a clear interpretation. In a risk–free model, the quantity $(\mathcal{W}_n - \mathcal{C}_n)$ vanishes, since we can, by assumption, perfectly reproduce the price–process of an option by hedging correctly. Therefore the correct hedge is independent of the time of sale of the contract. On the other hand, in a model with non–zero risk, one needs to correct the hedging strategy whenever the past hedge has not completely compensated the price change of the option. In particular, if the hedging strategy has made too much money and $\mathcal{W}_n > \mathcal{C}_n$has to decrease the amount of stock in the portfolio whenever the conditional drift μ_n is positive, and increase it otherwise. A similar result was obtained in [5], section 3.4; it was however noted there that this history dependent strategy, although reducing the variance, actually increases the probabilty of large losses.

8 Summary and conclusion

We have considered the problem of option pricing and hedging when stock returns are correlated in time. Using a variance minimization framework, we have obtained a general formula, valid for weakly correlated non–Gaussian processes. We have shown that in the limit of Gaussian price increments, the correlations are irrelevant, in the sense that the Black–Scholes formula holds with volatility that of the price increments on the time–scale of the re–hedging. For non–Gaussian processes, however, further non trivial corrections to the 'smile' effect are brought about by the correlations, even when the hedge is the Black–Scholes Δ–hedge. The kurtosis to be used in option–pricing is neither the local one, nor the global one, but some effective kurtosis which depends on the strength of the correlations. The above formalism can be extended to treat the case where the distribution of price increments depends explicitly on time, as would be needed to treat the case of a time dependent volatility. Finally, we have given formula for the optimal hedge when the expected return associated to options is non zero.

Appendix: Correlations within the Black–Scholes framework

The presence of correlations in the continuous time limit can be modelled as a history dependent drift term in the stochastic differential equation for the price (or log-price process). More precisely, the dynamical evolution of the price reads:

$$dX(t) = \mu(t)dt + \sigma_0 dW(t) \tag{17}$$

where the drift $\mu(t)$ depends on the whole past history and dW is a Brownian noise of unit variance per unit time. A simple model of correlations is to assume that $\mu(t)$

is constructed from the past values of $dW(t')$ as an exponential moving average:

$$\mu(t) = \epsilon\Gamma \int_{-\infty}^{t} dW(t') e^{-\Gamma(t-t')}. \tag{18}$$

When $\epsilon > 0$, this corresponds to correlations in price increments extending over a time scale Γ^{-1}. Conversely, $\epsilon < 0$ corresponds to anticorrelations. It is easy to see that in this model, the time dependent square volatility $\sigma^2(t)$ is given by:

$$\sigma^2(t) = \frac{1}{t} \int_0^t dt' \left(\sigma_0 + \epsilon(1 - e^{-\Gamma t - t'})\right)^2, \tag{19}$$

that interpolates between the short term volatility σ_0 and the long term volatility $|\sigma_0 + \epsilon|$.

In the Black–Scholes formalism, however, the drift is totally absent from the price and hedge of an option. This is intimately related to the fact that the Ito formula does not involve the drift $\mu(t)$ [7]. Hence, it is immediate that in this case, correlations are irrelevant and only the short term volatility σ_0 is needed to price the option. This is what we find within our general formalism in the case of Gaussian increments.

Note that the difference between the short term volatility and long term volatility is particularly striking in the case of an Ornstein-Uhlenbeck (mean reverting) process for which $\mu(t) = -KX(t)$. In this case, the long term volatility tends to zero: $X(t)$ has bounded fluctuations, even when $t \to \infty$. Correspondingly, the average pay-off over the objective probability distribution is finite for $t \to \infty$, whereas the option price, given by the standard Black–Scholes formula with a volatility σ_0, tends to infinity. The difference between the two comes from the hedging strategy that loses, because of the anticorrelations, an infinite amount of money in the limit $t \to \infty$.

References

[1] see *e.g.* : J.C. Hull *Futures, Options and Other Derivative Securities*, Prentice Hall (1997).

[2] P. Wilmott, *Derivatives, The theory and practice of financial engineering*, John Wiley (1998).

[3] J. P. Bouchaud & D. Sornette, *Journal de Physique I*, **4**, 863–881 (1994).

[4] M. Schweizer, *The Mathematics of Operations Research*, **20**, 1–32 (1995).

[5] J. P. Bouchaud & M. Potters, *Theory of Financial Risks*, Cambridge University Press (2000).

[6] see *e.g.* : J.Y. Campbell, A.W. Lo and A.C. MacKinlay *The Econometrics of Financial Markets*, Princeton (1997).

[7] see *e.g.* : M. Baxter, A. Rennie, *Financial Calculus*, Cambridge University Press (1996).