

Volatility and Dividends

Volatility Modelling with Cash Dividends and simple Credit Risk

Hans Buehler

Institut für Mathematik, MA 7-4
TU Berlin, Strasse des 17. Juni 136
10623 Berlin, Germany
e-mail: buehler@math.tu-berlin.de

Draft Version 0.94 December 28th, 2008

First draft March 2007

Updates at <http://www.math.tu-berlin.de/~buehler>

Abstract

This article revisits a simple and robust method on how to incorporate the handling of cash dividends and credit risk into equity derivatives pricing and risk management. In essence, we show that the stock price process upon default must have the form

$$S_t = (F_t - D_t)X_t + D_t$$

where X is a martingale with unit expectation, F is the forward curve and D is the floor imposed on the stock price process in the form of appropriately discounted future dividends.

We will show that the method presented is the only such method which is consistent with the assumption of cash dividends and simple credit risk, and we will demonstrate the impact on the pricing of various derivatives: plain Europeans, American options, Barriers and finally variance swaps and related derivatives.

This article is mainly based on results discussed in the first chapter of in “Equity Hybrid Derivatives”, Overhaus et al, Wiley 2006 [10] and the IQPC presentation “Options On Variance: Pricing And Hedging” given in 2006 [2].

1 Introduction

This article is about the very basics of equity modelling: the consistent incorporation of basic economic quantities into a simple, but efficient model of a share price process $S = (S_t)_{t \in \mathbb{R}_+}$ for the purpose of derivative pricing. We will focus on the following key ingredients into an asset’s price evolution:

- The underlying interest rates.
- Borrowing costs.
- Default risk.
- Dividends.
- The volatility structure.

We will take a pretty simplifying approach regarding most of these insofar as we assume that interest and borrowing rates, as well as default probabilities are deterministic and known in advance.

We will also use a particularly simple structure for the dividends of the stock: we assume that there is a series $0 = \tau_0 < \tau_1 < \tau_2 < \dots$ of dividend dates, at each of which a proportional and then a fixed cash dividend is paid. We denote the proportional dividends by β_j and the cash dividends by α_j . That means that at any dividend date the holder of one share receives dividends worth a cash value of $S_{\tau_j-}\beta_j + \alpha_j$; we use S_{τ_j-} to refer to the stock price just before the payment of the dividend.¹ Obviously, a dividend is only paid if the stock does not default up to the dividend date.

In the absence of tax imbalances, no-arbitrage arguments show that at any ex dividend date τ_j our share price process $S = (S_t)_t$ with a value of S_{τ_j-} just before the dividend date must drop according to

$$S_{\tau_j} = S_{\tau_j-}(1 - \beta_j) - \alpha_j . \quad (1)$$

In practise, the jump will occur at the opening of trading at the ex-dividend day. The non-trading period between the close and the open can either be modelled using zero volatility or by scaling time accordingly, i.e. by modelling the asset in its business time. The former has the advantage to allow more easily for multi-asset structures which are traded in different time zones.

2 Cash Dividends in Equity Modelling

The first step is now to derive a general structure of equity share price models which are consistent with the setting we introduced above. Note that we are not going to “approximate” a solution in some way. Rather, we derive the effects of introducing cash dividends for a share price process using purely no-arbitrage arguments. Technical details are provided in Overhaus et al. [10], chapter 1.

The deterministic instantaneous interest rate is denoted by $r = (r_t)_t$. That means that the price at time t of a risk-free zero coupon bond with maturity T is given as

$$P(t, T) = e^{-\int_0^T r_s ds} .$$

¹Note that we make the simplifying assumption that record date, ex-dividend date and payment date for the dividend all coincide.

Such a bond cannot default. In contrast, if the company issues a corporate zero bond (which will have zero recovery for the purpose of this article), an investor will demand a higher yield in compensation for the risk that the notional may actually not be paid.

Intuitively, in a market with diversifiable risk, we would expect the price of a risky bond to be equal to the price of a riskless bond times the (risk-neutral) probability of “survival” up to the maturity T of the instrument. Indeed, this is the case in our setup, which means that the value of a “risky” company bond provided it has not yet defaulted is given by

$$P^S(t, T) = P(t, T) \text{SV}(t, T) ,$$

where $\text{SV}(t, T)$ is the risk-neutral probability of survival up to T . At any point t , the function $T \mapsto \text{SV}(t, T)$ is usually extracted from Credit Default Swap (CDS) market prices. We write the survival probability function often as

$$\text{SV}(t, T) = e^{-\int_t^T h_s ds}$$

for some positive *credit spread* or *hazard rate* $h = (h_t)_t$. Since the default probabilities of the company for all maturities T are deterministic, h is also deterministic.

It is worthwhile noting that our inherent assumption by modelling default in this way is that we can not anticipate the default event itself by observing some publicly know variables (such as, say, the leverage of the company), but that the default is a sudden event (a fraud case, a sudden failure to honor a coupon payment on a bond etc). Mathematically, we say that the random default time τ is “inaccessible”, which means it cannot not be the time at which some publicly observable variable like the firm’s debt hits a barrier.

In general, it is far from obvious what exactly happens to a share price when such a default event occurs, but for sake of simplicity we will assume here that it drops to zero which makes our subsequent discussion particularly easy.

2.1 Consistent Share Price Processes

Assume that at some time $t > 0$, the stock has not yet defaulted and that we wish to enter into a forward contract with a maturity T . The standard way to compute the fair strike of the contract is by replication arguments: by re-investing all “relative” proceeds from holding the stock (repo and proportional dividends) to buy more shares, and by forward-selling the cash dividends we would receive from holding these shares, we can show that the the fair strike of the forward contract must be

$$F(t, T) = \left(R(t, T) S_t - \sum_{j: t < \tau_j \leq T} R(\tau_j, T) \alpha_j \right) \text{SV}(t, T) , \quad (2)$$

where we used the “proportional growth factor”

$$R(t, T) := e^{\int_t^T (r_s + h_s - \mu_s) ds} \prod_{j: t < \tau_j \leq T} (1 - \beta_j) .$$

The details of this derivation can be found in Overhaus et al. [10]. We will often abbreviate

$$F_t := F(0, t) \quad \text{and} \quad R_t := R(0, t) .$$

Note that the above formula (2) takes into account the possibility of a default after t and before T . One should note, in particular, that the value of the forward changes from the standard Black&Scholes result in the presence of cash dividends.

Given that a share bears no liability for the holder, the share price S_t can not become negative. This implies also that the forward can not be negative, because we would otherwise exchange at the maturity of the contract a negative cash amount for a share which will never have a negative value. That would clearly be an arbitrage situation. Hence, the forward is positive which in turn means by virtue of equation (2) that unless the company defaults, the value of the stock can never fall below a floor given by the “discounted” value of all remaining future cash dividends in the sense that

$$S_t \geq D_t \quad \text{where} \quad D_t := \sum_{j:\tau_j > t} \frac{\alpha_j}{R(t, \tau_j)} \quad \{\tau > t\} . \quad (3)$$

This is also economically sensible: the value of a share should always exceed the “discounted” expected value of all forthcoming dividends, in particular if these are tradable via dividend swaps (or calendar spreads of zero-strike calls). This means that the share price process $S = (S_t)_t$ is floored by the deterministic process $D = (D_t)_t$. While being economically sensible, this imposes a new way of looking at the mathematical modelling of the share price process: given that the stock price at time t cannot be below $D_t > 0$, it does not make sense to model it as a random process which can get arbitrarily close to zero, such as the Black&Scholes model,

$$S_t^{\text{BS}} := F_t \exp \left\{ \sigma(t) W_t - \frac{1}{2} \sigma(t)^2 t \right\} . \quad (4)$$

Instead, it makes much more sense to model the remaining randomness

$$S_t - D_t$$

with some process like the Black&Scholes model (4). Indeed, if we make sure the forward is preserved in our model, we get the following general summary of results:

RESULT 2.1 (Stock price with Cash Dividends and Credit) *All arbitrage-free stock price processes for S can be written as*

$$S_t = 1_{\tau > t} S_t^* \quad (5)$$

with “non-defaulting stock price”

$$S_t^* := (F_t^* - D_t) X_t + D_t \quad (6)$$

where the “pure stock price” $X = (X_t)_t$ is a non-negative (local) martingale with $X_0 = 1$.² In (6), we have used

$$F_t^* := \left(S_0 - \sum_{j: 0 < \tau_j \leq t} \alpha_j^* \right) R_t , \quad (7)$$

to denote the “risky forward” of the stock. It is the fair strike of a forward contract which settles only if the stock did not default until maturity. The formula is written in terms of the “discounted dividends”,

$$\alpha_j^* := \frac{\alpha_j}{R_{\tau_j}} .$$

Each α_j^* represent the time-zero value of all the cash dividends received per share bought at time 0 if all subsequent proportional proceeds from holding the share position are reinvested in the stock and if the stock did not default until the ex-dividend date τ_j .

Their value is consequently computed by “discounting” the future cash dividend α_j using the proportional riskless growth factor of the asset,

$$R_t := e^{\int_0^t (r_s + h_s - \mu_s) ds} \prod_{j: \tau_j \leq t} (1 - \beta_j) .$$

Since the value of the share at some future time t cannot drop below the “discounted” value of all forthcoming future dividends, the stock price cannot fallow below the floor $D = (D_t)_t$ given as

$$D_t := R_t \sum_{j: \tau_j > t} \alpha_j^* .$$

This formula simply discounts the remaining dividends consistently to t , cf. (3).

Also note that (5) readily implies that

$$F_t = F_t^* \text{SV}(0, t) .$$

Interpretation: The essence of the preceding result is that we can separate the modelling of the *volatility risk* of the stock, X , entirely from the other characteristics of the equity. In fact,

$$S_t^* = (F_t^* - D_t)X_t + D_t$$

really means that the only thing left to do is to specify the “pure” stock process X (hence the name). In this context, note that the process X itself should not carry any default risk (i.e., reach zero) because it would not actually imply default on the outstanding dividends: if $X_t = 0$, then the martingale property of X forces it to stay in 0, such that the stock S reduces to a pure risky cash bond paying the dividends as coupons.

²Here, we will always assume that X is a true martingale. Also note that strictly speaking, the “all” above is only true if we assume “no free lunch with vanishing risk”, (cf. [4]) which is stronger than “no arbitrage”.

REMARK 2.1 (Business time) *To ease the implementation of business time such that the stock has zero volatility during non-trading hours, it is convenient to rescale time for X , i.e. write it in terms of a suitable martingale Z as*

$$X_t := Z_{A_t}$$

where $(A_t)_t$ is the business time process for the stock price. In particular, A does not increase outside trading hours. Moreover, it can also be used to model “event dates” where time is faster in the sense that the daily volatility is a multiple of a standard day (earnings announcements, index rebalancing, etc).

Stochastic Dynamics

Another convenient notation for the dynamics of the stock price is derived as follows: let us denote by

$$S_0^* := S_0 - \sum_{j=1}^{\infty} \alpha_j^* , \quad (8)$$

the “effective” stock price level which indicates the part of the initial stock price S_0 which is actually subject to volatility risk. Then, we obtain from (6) the key formula

$$S_t^* := S_0^* R_t X_t + D_t , \quad (9)$$

which also implies $F_t^* = S_0^* R_t + D_t$. Moreover, it allows us to apply Ito (cf. section A) which yields³ the SDE for S^* ,

$$\begin{aligned} dS_t^* &= S_{t-}^* (r_t + h_t - \mu_t) dt \\ &\quad + (S_{t-}^* - D_{t-}) \frac{dX_t}{X_{t-}} \\ &\quad - \sum_j \{ S_{t-}^* \beta_j + \alpha_j \} \delta_{\tau_j}(dt) . \end{aligned} \quad (10)$$

This equation separates the dynamics of the non-defaulting stock price S^* into its three main components: the first line represents the continuous growth of the stock, the second line its “randomness” (note that the randomness only applies to the excess of the stock level over the future dividends), and the last line represents the impact of dividend payments.

The SDE for S is then

$$dS_t = 1_{\tau > t} dS_t^* - S_{t-}^* \delta_{\tau}(dt) . \quad (11)$$

³The continuous differentials of R and D , respectively, are given as $dR_t^c = R_{t-}(r_t + h_t - \mu_t) dt$ and $dD_t^c = D_{t-}(r_t + h_t - \mu_t) dt$. Applying Ito to (9) hence gives

$$\begin{aligned} dS_t^* &= S_0^* X_{t-} dR_t^c + S_0^* R_{t-} X_{t-} \frac{dX_t}{X_{t-}} + dD_t^c + \sum_j \{ S_{t-}^* - S_{t-}^* \} \delta_{\tau_j}(dt) \\ &= S_{t-}^* (r_t + h_t - \mu_t) dt + (S_{t-}^* - D_{t-}) \frac{dX_t}{X_{t-}} + \sum_j \{ S_{t-}^* (1 - \beta_j) - \alpha_j - S_{t-}^* \} \delta_{\tau_j}(dt) , \end{aligned}$$

which yields the assertion after rearrangement (we implicitly assumed that $X > 0$ and that the jumps of X do a.s. not coincide with the dividend dates).

The Total Return

The “total return” $S^{(\text{TR})}$ of a stock S (most commonly an index) is generated by re-investing all dividends (but not any proceeds from the repo) back into the asset. That simply means that the value of the total return index will not drop at a dividend date. Therefore, the total return must have a form

$$S_t^{(\text{TR})} = S_0 e^{\int_0^t (r_s + h_s - \mu_s) ds} Y_t \quad \{\tau > t\} \quad (12)$$

where Y is a martingale. If Y is continuous, then the total return is essentially the “continuous part” of the stock price. Mathematically, this can be written as

$$S_t^{(\text{TR})} = \int_0^t dS_{u-} \quad \{\tau > t\} .$$

More generally, if t is between the dividend dates τ_k and τ_{k+1} , we have

$$S_t^{(\text{TR})} = \frac{S_t}{S_{\tau_k}} \prod_{j=1}^k \frac{S_{\tau_j-}}{S_{\tau_{j-1}}} \quad \{\tau > t\} .$$

We have therefore identified the martingale Y in equation (12) as

$$Y_t = \frac{S_t}{S_{\tau_k} R(\tau_k, t)}$$

whenever $\tau_k \leq t < \tau_{k+1}$. In particular, Y is *not* equal to X .

Hence, $S^{(\text{TR})}$ is *not* Markov with respect to itself: we need to simulate S in order to be able to compute a payoff depending on the total return.

REMARK 2.2 *Note that the total return is also bounded from below away from zero. So it cannot be modelled using Black&Scholes or a similar process, either.*

2.2 European Option pricing

After we have identified with equation (6) a formulation of the share price process which is consistent without our dividend assumptions, the question is which impact this has on the pricing of options.

We first focus on European options and the implications therefrom. The first most obvious observation is that a European option on S is essentially a European option on X , since we can write for each payoff H

$$H(S_T) = 1_{\tau > T} H\left((F_T^* - D_T) X_T + D_T\right) + 1_{\tau \leq T} H(0)$$

Now assume that we can observe market prices for call options on S

$$C(T, K) := P(0, T) \mathbb{E} \left[(S_T - K)^+ \right]$$

for all strikes $K \geq 0$ and maturities $T \geq 0$. Let us also denote by

$$\mathbb{C}(T, x) := \mathbb{E} \left[(X_T - x)^+ \right] ,$$

a call on the pure stock price X with “relative strike” x .

RESULT 2.2 *The price of a call on the pure stock price X can be expressed in terms a call on S as*

$$\mathbb{C}(T, x) = \frac{1}{P^S(0, T)(F_T^* - D_T)} C\left(T, (F_T^* - D_T)x + D_T\right). \quad (13)$$

Reversely,

$$C(T, K) = P^S(0, T)(F_T^* - D_T) \mathbb{C}\left(T, \frac{K - D_T}{F_T^* - D_T}\right). \quad (14)$$

Implied Local Volatility

In particular, this means that the knowledge of all option prices for S for each maturity T gives us access to the marginal densities of the underlying pure martingale X via

$$\mathbb{P}[X_T = x] = \partial_{xx}^2 \mathbb{C}(T, x).$$

Furthermore, Dupire's implied local volatility [6] for the process X is given as⁴

$$\varsigma_t^X(x)^2 := \frac{\partial_t \mathbb{C}(T, x)}{2 x^2 \partial_{xx}^2 \mathbb{C}(T, x)}. \quad (15)$$

Using a local volatility for X rather than for S is much less prone to numerical instabilities due to dividend or credit effects, since all these effects have essentially been stripped out. For example, (10) implies that a local volatility function for S directly would not be properly defined below the (time-dependent) barrier D . Moreover, this formulation in conjunction with remark 2.1 allows the calibration of business-time local volatility.

REMARK 2.3 *Of course, with $dX_t/X_t = \varsigma_t(X_t)dW_t$, we can rewrite (10) as*

$$dS_t = S_{t-}(r_t + h_t - \mu_t) dt + S_{t-}\sigma_t(S_{t-}) dW_t - \sum_j \{S_{t-}^* \beta_j + \alpha_j\} \delta_{\tau_j}(dt)$$

by setting

$$\sigma_t(s) := \begin{cases} \frac{s - D_{t-}}{s} \varsigma_t\left(\frac{s - D_{t-}}{F_{t-}^* - D_{t-}}\right) & (s > D_{t-}) \\ 0 & (s \leq D_{t-}) \end{cases}. \quad (16)$$

Note that the left limits on D and F^ indicate that these quantities shall not contain the dividends paid at t .*

⁴That means that if a unique solution $\tilde{X}$ to $d\tilde{X}_t/\tilde{X}_t = \varsigma_t^X(\tilde{X}_t)dW_t$ with $\tilde{X}_0 = 1$ exists, then it reprices the market, $\mathbb{E}[(\tilde{X}_T - x)^+] = \mathbb{C}(T, x)$.

No-Arbitrage with Cash Dividends

This is all well, but further investigation of equation (13) highlights a potential issue: we know that the imposition of cash dividends demands that the stock price S_T cannot fall below the floor D_T unless it defaults. Hence, there is a limit on the value of the call on the stock with a strike K of less than D_T : such a call must be equal to the discounted forward value, $C(T, K) = P(0, T)(F_T - K) = P^S(0, T)(F_T^* - K)$. This is a hard constraint on the value of calls with strikes below K which is not present when we use purely proportional dividends as in Black&Scholes.

In fact, it should come to no surprise that the introduction of cash dividends and credit risk changes the usual “no-arbitrage”-conditions. Consider, as another example, a call calendar spread over a dividend date: a call just after a dividend date is worth less than a call which is struck with the same strike just before the dividend. That means that the call prices are *not* always increasing in time.

The common no-arbitrage conditions such as positivity of butterflies and calendar spreads plus some constraints on the boundaries of course still hold, but not for S itself, but for the pure stock price X . They are summarized here:

RESULT 2.3 *Assume calls $\mathbb{C}(T, x)$ are given for all maturities $T \geq 0$ and all strikes $x \geq 0$. Then, X is a strictly positive true martingale if and only if*

1. *The forward is preserved, $\mathbb{C}(T, 0) = 1$. This implies that $\mathbb{C}(T, 0) = F_T$.*
2. *Butterflies have a positive value, $\partial_{xx}^2 \mathbb{C}(T, x) \geq 0$.*
3. *Calendar spreads are positive, $\partial_T \mathbb{C}(T, x) \geq 0$.*
4. *The probability of reaching zero is nil, i.e. $\partial_x \mathbb{C}(T, 0) = -1$.⁵*
5. *For high strikes we have $\lim_{x \rightarrow \infty} \mathbb{C}(T, x) = 0$.*

In this case we say that the market is “strictly free of arbitrage”.⁶

The above summary is not only useful in the context of cash dividends; it also tells us that the slope of the call prices as a function of strike must change in the presence of credit risk. In particular, assume that there are no dividends but that there is a substantial credit risk, i.e. $SV(0, T) \ll 1$. In this case,

$$\frac{\partial_K C(T, 0)}{P(0, T)} = -SV(0, T) \gg -1 ,$$

hence the (undiscounted) call prices as a function of K must have a much higher slope than -1. Such a distribution cannot be generated by a diffusion. This

⁵One could weaken this condition to $\partial_x \mathbb{C}(T, 0) \geq -1$, in which case X can reach zero (for example, in the CEV model). However, as described above, we consider a stock price model in which the stock can degenerate into a coupon-bearing risky bond as not desirable.

⁶The case where only a discrete set of option prices is observed is equally straight-forward, see [3].

means in particular that the standard measure of volatility risk, Black&Scholes' implied volatility, will explode as $K \downarrow 0$.⁷ To recall, the (*standard*) *implied volatility of S* for strike K and maturity T is the positive number $\sigma_{\text{BS}}(T, K)$ which solves

$$\mathbb{BS}\left(\sigma_{\text{BS}}(T, K), T, \frac{K}{F_T}\right) \equiv \frac{C(T, K)}{P(0, T)F_T} . \quad (17)$$

where

$$\mathbb{BS}(\sigma, T, k) := \mathcal{N}(d_+) - k\mathcal{N}(d_-) \quad \text{with} \quad d_{\pm} = \frac{-\ln k \pm \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} .$$

In fact, the presence of even a minimal credit risk imposes very extreme lower bounds on implied volatility. Figure 1 shows a numerical example of an implied volatility surface when there is an 0.0001% percent probability of default per year.

Figure 1: Black&Scholes implied volatilities for a hypothetical stock price with a CDS spread of 5 and a flat pure stock price volatility of X .

Implied Volatilities

In other words, *standard* implied volatility in its simple form may not be the best way of seeing the pure “volatility risk” in the stock price. Indeed, once we have accepted thinking of X as the volatile part of the stock price, it is much more natural to look at the implied volatilities of X directly. Thanks to equation (13) this is straightforward:

DEFINITION 2.1 *We call $\sigma_X(T, x)$ the “pure” implied volatility of S (or X) if it solves*

$$\mathbb{BS}(\sigma_X(T, x), T, x) \equiv \mathbb{C}(T, x) .$$

Note that this means that

$$C(T, K) = P^S(0, T) (F_T^* - D_T) \mathbb{BS}\left(\sigma_X\left(T, \frac{K - B_T}{F_T^* - D_T}\right), T, \frac{K - B_T}{F_T^* - D_T}\right) . \quad (18)$$

If our new measure of implied volatility is used, then it is possible to mark the implied volatility of X rather freely – for example, by using an implied volatility interpolation such as Gatheral's [7] or by specifying some martingale-dynamics such as Heston's for X – without violating the inherent arbitrage bounds on the standard implied volatility surface of S which are implied by the presence of

⁷See also [8] where it is shown that the asymptotic growth rate of implied volatility is of at most of degree 2 in log-strike.

credit risk and dividends. In particular, such a surface prevents trivial arbitrage between equity and credit instruments.

As an example, figure 2 compares the implied volatilities of .FTSE100 with cash dividends vs. purely proportional dividends. The consequence of impos-

Figure 2: A fitted implied volatility surface for .FTSE100 under the assumption cash dividends for the next eight years. We see both the pure and the resulting standard implied volatility surface. Note the slight downward slope of the implied volatility skew.

ing cash dividends results in a higher pure implied volatility compared to the pure proportional case, in particular for lower strikes. This can be explained by considering an option with a strike K just above the floor D_T . While such an option will have a small but non-zero optionality value⁸ for the purely proportional case for most volatilities, it will have virtually no optionality in the presence of cash dividends unless the pure volatility of X is very very high. Figure 3 shows the relation between “cash volatility” and “proportional volatility” in more detail. The reverse conclusion is that if we mark the implied volatility

Figure 3: The pure implied volatility curve given a standard implied volatility curve.

of an underlying using its pure volatility, then the resulting standard implied volatility will actually be lower for low strikes, and it will become zero below the floor D_T .

While the introduction of cash dividends *lowers* the standard implied volatility, the opposite effect occurs, though, if we start to look at single stocks which carry some credit risk. In this case, the relationship between standard and pure implied volatility is reversed: even a small probability of default puts a non-zero probability point-mass at the default state zero, which lets standard implied volatilities explode. This is illustrated in figure 4 where we show the standard and pure implied volatility of CREDITRISK. Hence, even a small risk of default immediately produces the extreme skews observed for many single name implied volatility surfaces.

Figure 4: A fitted implied volatility surface for CREDITSTICK which has a CDS spread of XYZ. Note the extreme impact of the credit risk on the implied volatility surface.

⁸excess of option value above intrinsic

2.3 PDE Pricing

To price callable or auto-callable options on just a few underlyings, it is often recommendable to use PDEs instead of Monte-Carlo. We will briefly address the issue of pricing with backward schemes. For simplicity of exposure, let us first assume that X is a local-volatility process, i.e.

$$\frac{dX_t}{X_t} = \varsigma_t(X_t) dW_t \quad (19)$$

as suggested in (15) above where ς is sufficiently well-behaved such that $X > 0$.

Let $H(S_T)$ be a payoff at maturity and let us assume first that there is no immediate recovery values for the option, i.e. the holder simply receives $H(0)$ at maturity if the stock defaulted. Let

$$H_t(s) := P(t, T) \mathbb{E} [H(S_T) | S_t = s] \quad (20)$$

be the value of the claim at any time t . Using (10) we get

$$\begin{aligned} dH_t(S_t) &= dH_t(1_{\tau > t} S_t^*) \\ &= \left\{ H_t(0) - H_t(S_{t-}) \right\} \delta_\tau(dt) \\ &\quad + \partial_s H_t(S_{t-}) \left\{ S_{t-}(r_t + h_t - \mu_t) dt + (S_{t-} - 1_{\tau \geq t} D_{t-}) \varsigma_t(X_t) dW_t \right\} \\ &\quad + \frac{1}{2} \partial_{ss}^2 H_t(S_{t-}) \left((S_{t-} - 1_{\tau \geq t} D_{t-}) \varsigma_t(X_t) \right)^2 dt \\ &\quad + \sum_j \left\{ H_t(S_{t-}(1 - \beta_j) - \alpha_j 1_{\tau \geq t}) - H_t(S_{t-}) \right\} \delta_{\tau_j}(dt) \\ &\quad + \partial_t H_t(S_{t-}) dt \end{aligned}$$

By construction, this must be a discounted martingale,⁹ hence we can derive the associated PDE from the above SDE. Here, it is convenient to recall the volatility function

$$\sigma_t(s) := \begin{cases} \frac{s - D_{t-}}{s} \varsigma_t\left(\frac{s - D_{t-}}{F_{t-}^* - D_{t-}}\right) & (s > D_{t-}) \\ 0 & (s \leq D_{t-}) \end{cases},$$

as introduced in (16) above. Also note that the sets $\{\tau \geq t\}$ and $\{S_{t-} > 0\}$ coincide given our requirement that $X > 0$. This allows us to write

$$\begin{aligned} 0 &= \partial_t H_t(s) dt - H_t(s) r_t + \partial_s H_t(s) s (r_t + h_t - \mu_t) + \frac{1}{2} \partial_{ss}^2 H_t(s) s^2 \sigma_t(s)^2 \\ &\quad + 1_{s > 0} \sum_j \left\{ H_t(s(1 - \beta_j) - \alpha_j) - H_t(s) \right\} \delta_{\tau_j}(dt) \\ &\quad + (H_t(0) - H_t(s)) h_t 1_{s > 0}. \end{aligned}$$

⁹To this end, note that $N_t := 1_{\tau > t} + \int_0^t \hat{h}_s^\tau ds$ is a martingale: assume P is a Poisson-process with intensity h . Then, its compensated version $N'_t := P_t - \int_0^t h_s ds$ is a martingale. Now denote by τ' its first jump time, i.e. $\tau' := \inf\{t : N'_t = 1\}$ and define N'' as the stopped process N' . Then, N'' and $1 - N$ have the same distribution.

This equation offers itself to straight forward interpretation of the dynamics of the stock price: the first line is simply the standard PDE for a local volatility-driven stock price with a modified drift which reflects the additional average growth rate which is to compensate the stock holder for the additional risk of default. The second line is the jump condition for the dividend payments.

The last line, finally, represents the infinitesimal probability of defaulting, in which case the value drops to $H_t(0)$. Note that the multiplication with the indicator $1_{s>0}$ is redundant in the case here, but it becomes necessary if we extend the above formula to take into account “immediate recovery” values: this simply means that instead of falling to $H_t(0)$, the value of the security becomes, say, ρ_t . This quantity could technically also depend on the previous stock price level.

A good example of a security whose pricing requires such an additional recovery term ρ is a convertible bond which will default to the recovery value of its intrinsic bond, which is usually assumed to be around 40% of notional.

RESULT 2.4 *The PDE for an option with terminal payoff H and recovery process ρ is*

$$\begin{aligned} 0 = & \partial_t H_t(s) dt - H_t(s) r_t + \partial_s H_t(s) s (r_t + h_t - \mu_t) + \frac{1}{2} \partial_{ss}^2 H_t(s) s^2 \sigma_t(s)^2 \\ & + 1_{s>0} \sum_j \left\{ H_t(s(1 - \beta_j) - \alpha_j) - H_t(s) \right\} \delta_{\tau_j}(dt) \\ & + (\rho_t - H_t(s)) h_t 1_{s>0} . \end{aligned} \quad (21)$$

It should be noted that this equation also holds for the dynamics of the price process of the security after default: once default occurred, the value process of the option is trivial and has to accrue at the risk-free rate.

American Options

American options are priced as usual: assume that upon exercise at time t , the option holder receives $E_t(S_{t-})$. Accordingly, the respective partial differential inequality is given by

$$\begin{aligned} E_t(s) \geq & \partial_t H_t(s) dt - H_t(s) r_t + \partial_s H_t(s) s (r_t + h_t - \mu_t) + \frac{1}{2} \partial_{ss}^2 H_t(s) s^2 \sigma_t(s)^2 \\ & + 1_{s>0} \sum_j \left\{ H_t(s(1 - \beta_j) - \alpha_j) - H_t(s) \right\} \delta_{\tau_j}(dt) \\ & + (\rho_t - H_t(s)) h_t 1_{s>0} . \end{aligned}$$

PDE for the pure stock price process

The above PDE (21) is the natural description of the asset dynamics, but when implemented numerically, it has two major drawbacks: first is the need to implement a jump condition on the finite difference grid at every dividend date which will slow down convergence in particular when used with indices which

may pay dividends nearly every day such as S&P500. Secondly, (21) requires at least in its pure application the evaluation of unnatural “instantaneous” quantities such as r_t and h_t . Unless we actually require the value of the option at any time between now and maturity (continuous exercise conditions or continuous Barriers), it is therefore often more efficient to express the price of an option on S as a price of an option on X .

To this end, consider a Bermudan option that can be exercised on dates $0 = T_0 < \dots < T_N$ with exercise values $E_0(s), \dots, E_N(s)$ depending on the stock price s at that time. We assume that the exercise decision is communicated after any dividends are paid on the stock. In the event of default, the holder of the option recovers ρ_t .

At maturity, assuming no default, the value of the option is obviously

$$H_{T_N}(s) = \max \{E_N(s), 0\} .$$

At any exercise date T_ℓ before maturity and default, standard pricing theory tells us that the value of the option will be the maximum between exercising the option now or holding on to it. If we hold on to the option, the stock may default before the next exercise decision. Hence, the holding value must be the sum of the discounted value of the next exercise date, multiplied by the survival probability, plus the value of the recovery in the event of default, weighted by the default probability. We write this as

$$H_{T_\ell}(s) := \max \left\{ E_\ell(s), C_\ell(s) + R_\ell \right\} ,$$

where

$$C_\ell(s) := P^S(T_\ell, T_{\ell+1}) \mathbb{E} \left[H_{T_{\ell+1}}(S_{T_{\ell+1}}) \mid S_{T_\ell} = s \right] \quad (22)$$

denotes the value of holding the option if it does not default until the next exercise time, and where

$$R_\ell := \mathbb{E} \left[P(T_\ell, \tau) \rho_\tau 1_{T_\ell < \tau \leq T_{\ell+1}} \mid \tau > T_\ell \right] = \int_{T_\ell}^{T_{\ell+1}} h_t \rho_t P^S(T_\ell, t) dt$$

denotes the recovery value if it does. Note that R_ℓ is deterministic in our current framework. To compute (22), define for $t \in [T_\ell, T_{\ell+1}]$

$$\tilde{C}_t(x) := \mathbb{E} \left[H_{T_{\ell+1}} \left((F_{T_{\ell+1}}^* - D_{T_{\ell+1}}) X_{T_{\ell+1}} + D_{T_{\ell+1}} \right) \mid X_t = x \right] \quad (23)$$

such that

$$C_\ell(s) = P^S(T_\ell, T_{\ell+1}) \tilde{C}_{T_\ell} \left(\frac{s - D_{T_\ell}}{F_{T_\ell}^* - D_{T_\ell}} \right) .$$

Under the local volatility assumption (19) for X , the PDE for (23) is as usual given by

$$0 = \partial_t \tilde{C}_t(x) dt + \frac{1}{2} \partial_{xx}^2 \tilde{C}_t(x) x^2 \varsigma_t(x)^2 ,$$

which is much easier to implement efficiently, in particular if ς is given analytically via Dupire’s formula.

2.4 Monte-Carlo

From an implementation point of view, Monte-Carlo methods are often more intuitive even if they suffer from less efficiency for low dimensions. However, for dimensions above two or for strongly path-dependent products, there is virtually no alternative so far for the evaluation of exotic options.

In our framework, the implementation of an efficient Monte-Carlo scheme can be achieved by separating the simulation of X which depends on the choice of the stochastic dynamics and the subsequent application of dividend and credit effects. In conjunction with remark 2.1 on business time, such an implementation can be transparent from a modelling point of view: the handling of X does not require any additional information on S . For example, the transformations described in result 2.2 can be used to provide a “pure” implied volatility surface for the calibration of the stochastic process X .

3 Variance Swaps

While the previous concepts are straight-forward for simple European options – after all, we simply apply an affine transformation to the stock price –, the relationship between S and X of course also impacts the pricing and risk management of variance swaps.

The basic idea of a variance swap is to pay the annualized realized variance, i.e. the squared volatility, of the returns of the underlying asset in exchange for a previously agreed fixed “set volatility”, also squared. The definition of realized variance over close-of business days $0 = t_0 < \dots < t_n = T$ is typically simply the standard mathematical definition,

$$RV^I(T) := \sum_{i=1}^n \left(\log \frac{S_{t_i}}{S_{t_{i-1}}} \right)^2 .$$

However, to avoid that dividends impact the computation of realized variance, single-stock variance swaps are usually based on a dividend-adjusted version,

$$RV^S(T) := \sum_{i=1}^n \left(\log \frac{S_{t_i}^{(TR)}}{S_{t_{i-1}}^{(TR)}} \right)^2 ,$$

which is just a clever way of saying that the returns of S as adjusted for the dividends.¹⁰

With either definition of realized variance, the payoff of a variance swap is

$$\frac{252}{n} RV(T) - K^2$$

where K represents the fixed leg’s “set volatility”. The factor $252/n$ annualizes the variance assuming that the year has 252 business days. The number 252 is usually contractually fixed.

¹⁰ Assuming that dividend effects occur end-of day, $RV^S(T) = \sum_{i=1}^n (\log S_{t_i-}/S_{t_{i-1}})^2$.

REMARK 3.1 *The payoff of single-stock variance swaps is usually also subject to an additional cap at 150% of K^2 since very extreme stock price movements, for example due to takeover rumors or default, would otherwise have a very strong effect on the realized variance estimator. This turns the contract into a non-linear payoff on realized variance.*

For the time being, we will ignore the cap.

For pricing purposes, it should be clear that the fixed leg of a variance swap is a trivial zero bond. Moreover, the scaling factor $252/n$ in front of realized variance simply changes the notional. Hence, when talking about “variance swaps”, we will for most part assume that K is zero (a so-called “zero-strike variance swap”) and that the scaling factor is 1. In other words, a variance swap henceforth simply pays realized variance to the holder unless otherwise stated.

3.1 Classic Variance Swap Pricing

The classic approach to price a (zero-strike, not annualized) variance swap is to assume that X has no jumps, that no cash dividends are paid, the repo-rate is zero and that there is no risk of default. Then, Ito’s formula applied to the logarithm of the stock for each interval $[t_{i-1}, t_i]$ yields

$$\begin{aligned} \log \frac{S_{t_i}}{S_{t_{i-1}}} &= \int_{t_{i-1}}^{t_i} \frac{dS_t}{S_t} - \frac{1}{2} \int_{t_{i-1}}^{t_i} \frac{d\langle S \rangle_t}{S_t^2} \\ &\approx \frac{1}{S_{t_{i-1}}} (S_{t_i} - S_{t_{i-1}}) - \frac{1}{2} \frac{1}{S_{t_{i-1}}^2} (S_{t_i} - S_{t_{i-1}})^2 . \end{aligned} \quad (24)$$

Summing these terms up and rearranging them gives Neuberger’s [9] famous formula

$$RV^I(T) \approx -2 \log \frac{S_T}{S_0} + 2 \sum_{i=1}^n \frac{1}{S_{t_{i-1}}} (S_{t_i} - S_{t_{i-1}}) . \quad (25)$$

This formula suggests the following approach: to replicate the left hand side, short two European log-contracts (we comment on this below) and execute a daily delta-hedge which results in the last term in (25), following the remarks in [10]. This means running a delta of

$$\Delta_i^{\text{hedge}} := 2 \frac{P(t_i, T)}{S_{t_{i-1}}} , \quad (26)$$

which amounts to a “cash delta” of 2, properly discounted (“cash delta” denotes the amount of capital invested in the stock, i.e. it is the standard delta times the stock price). Note, however, that this strategy is not self-financing if rates are not zero. Its costs are covered by the discounted value of the last term in equation (25).¹¹

¹¹At the beginning of each period, we borrow a cash delta amount of $2P(t_i, T)$ and invest it into the stock for a price of $S_{t_{i-1}}$. At the end of the period, we sell the stock position for S_{t_i} ,

Intra-day Pricing and Risk Management

A theoretically slightly more accurate approach than using (25) is to use a continuous form of that expression. The basic idea is that RV^I is an unbiased estimator for the quadratic variation of the log-returns of the stock, $QV(t) := \langle \log S \rangle_t$ even if S has jumps. This means that

$$\mathbb{E} [RV^I(T) | \mathcal{F}_{t_i}] = \mathbb{E} [QV(T) | \mathcal{F}_{t_i}] \quad (27)$$

for all $i = 0, \dots, N$. We also set $QV(T_1, T_2) := QV(T_2) - QV(T_1)$. Now, we simply apply Ito to the log-contract and rearrange accordingly. This gives:

$$QV(t, T) = -2 \log \frac{S_T}{S_t} + 2 \int_t^T \frac{dS_u}{S_u}, \quad (28)$$

which for the case $t = 0$ essentially coincides under expectation with (25). Note that in the absence of credit risk and cash dividends, (10) shows that

$$QV(t, T) = -2 \log \frac{S_T}{F(t, T)} + 2 \int_t^T \frac{dX_u}{X_u}. \quad (29)$$

To finally evaluate a variance swap at a point in time $t \in [t_{i-1}, t_i]$, we split the computation of the expected payout into three distinct parts:

$$\begin{aligned} & \mathbb{E} [RV^I(T) | \mathcal{F}_t] \\ &= \sum_{k=1}^{i-1} \left(\log \frac{S_{t_k}}{S_{t_{k-1}}} \right)^2 + \mathbb{E} \left[\left(\log \frac{S_{t_i}}{S_{t_{i-1}}} \right)^2 | \mathcal{F}_t \right] + \mathbb{E} \left[\sum_{k=i+1}^n \left(\log \frac{S_{t_k}}{S_{t_{k-1}}} \right)^2 | \mathcal{F}_t \right] \\ &= \underbrace{RV^I(t_{i-1})}_{\text{Past}} + \underbrace{\mathbb{E} \left[\left(\log \frac{S_{t_i}}{S_{t_{i-1}}} \right)^2 | \mathcal{F}_t \right]}_{\text{Present}} + \underbrace{\mathbb{E} [QV(t_i, T) | \mathcal{F}_t]}_{\text{Future}} \end{aligned} \quad (30)$$

While the “Past” is trivial to compute in theory (apart from practical problems such as stock splits, non-trading days and other real-life phenomena), the question remains how to compute “Present” and “Future”. Since the main idea of splitting the payoff is to guarantee that the pricing scheme returns the correct cash delta, it is in fact sufficient to compute the “Present” with a zero-volatility Black&Scholes model, in which case the “Future” is in the light of (29) simply given by

$$P(t, T) \mathbb{E} [QV(t_i, T) | \mathcal{F}_t] = P(t, T) \mathbb{E} \left[\log \frac{S_T}{F(t, T)} | \mathcal{F}_t \right].$$

(This formula is also true if purely proportional dividends are present.)

hence our P&L is $2P(t_i, T)S_{t_i}/S_{t_{i-1}} - 2P(t_i, T)/P(t_{i-1}, t_i)$, which translates at maturity into

$$2(S_{t_i}/S_{t_{i-1}} - 1) - 2(1/P(t_{i-1}, t_i) - 1) = 2(S_{t_i}/S_{t_{i-1}} - 1) - 2(F_{t_i}/F_{t_{i-1}} - 1),$$

where the latter term matches exactly the expectation of the last term in (25).

Realized Variance vs. Quadratic Variation

The approximation of realized variance by quadratic variation is valid only for linear payoffs on variance. The price of any non-linear payoff such as a call on realized variance must be computed based on realized variance. To see this, assume that X is a Black&Scholes process and that $S = X$. Following Christopher Jordinson's argument presented on page 12ff. in [2] we then have

$$\frac{RV(0, T)}{QV(0, T)} = \sum_{i=1}^n \left\{ \frac{x_i - \frac{1}{2}\sigma\sqrt{T/n}}{\sqrt{n}} \right\}^2$$

with $(x_i)_{i=1, \dots, n}$ iid standard normal. Consequently, the ratio RV/QV is chi-square distributed with parameter $\lambda = \sigma^2 T / 4n^2$. Figure 5 illustrates the effect on option prices in a Heston model: on the short end the value of an option written on quadratic variation declines while it increases for options on realized variance.




Figure 5: Options on realized variance vs. options on quadratic variation in a Heston model. The graph illustrates that short term options on realized variance cannot be priced using quadratic variation.

The Delta of a Variance Swap in Black&Scholes is Zero

From (30), we can easily derive the delta of a variance swap as

$$\Delta_t^{\text{VS}} = 2 \frac{P(t, T)}{S_t} \log \frac{S_t}{S_{t_{i-1}}} + \partial_{S_t} P(t, T) \underbrace{\mathbb{E} \left[\log \frac{S_T}{F(t, T)} \mid \mathcal{F}_t \right]}_{(*)}. \quad (31)$$

Here, it is important to notice that the term $\log S_T/F(t, T)$ does not depend on the current stock price level, S_t . That means any sensitivity of $(*)$ towards S_t must be via a dependency of the volatility surface on the stock price level, as it is the case for local volatility models. However, for all ‘classic stochastic volatility models, jump models or, in fact, Black&Scholes, where the dynamics of the underlying volatility structure do not functionally depend on S_t (or X_t for that matter), the term $(*)$ will have no delta.

REMARK 3.2 *The delta of a variance swap in Black&Scholes is zero,*

$$\Delta_{t_i}^{\text{VS}} = 0,$$

at any fixing date $t = t_i$.

The same is true for all classic stochastic volatility and jump models.

The famous delta-term (26) which we quoted above is the result of offsetting a variance swap position with the respective log-contract: assume that we add a log-contract initiated at time $t = 0$ to the payoff (30). This contract is

$$\mathbb{E} \left[\log \frac{S_T}{F(0, T)} \mid \mathcal{F}_t \right] = \mathbb{E} \left[\log \frac{S_T}{F(t, T)} \mid \mathcal{F}_t \right] + \log \frac{S_t}{S_0} - \log P(0, t)$$

where we used the fact that $F(t, T)/F(0, T) = S_t/S_0/P(0, t)$ in our current setting. If we add the delta of this term to (31) we end up at any fixing date t_i with exactly the previously mentioned

$$\Delta_i^{\text{hedge}} = 2 \frac{P(t_i, T)}{S_{t_i}}.$$

From a booking perspective, this implies that the variance swap and its log-contract hedge (or any approximation to it) should be booked into the same account in order to make sure that the system reports the right combined stock sensitivity. Note that such a procedure also allows to book much reduced “log contracts” against the variance swap, such as simple butterflies, while the price of the variance swap is still computed using the full log-contract.

REMARK 3.3 *If the log-contract $(*)$ is computed using the approximation approach discussed below, then a typical risk-management system will produce a delta for this term based on the rules governing the change of implied volatility if the stock price changes. In this sense, the delta of a variance swap in such a system is a pure “dVol/dSpot” measure.*

Log-contracts

The above discussion and in particular the replication formula (25) assumes that we can actually trade log-contracts. In practise, this is not the case. If we were able to trade arbitrary amounts of European options with all strikes, we could replicate any European payoff by virtue of the following formula:

$$f(x) - f(x_0) = f'(x_0)(x - x_0) + \int_{x_0}^{\infty} f''(z)(x - z)^+ dz + \int_0^{x_0} f''(z)(z - x)^+ dy dz . \quad (32)$$

For the log-contract, we obtain:

$$\log \frac{S_T}{F_T} = \frac{S_T - F_T}{F_T} + \int_{F_T}^{\infty} \frac{(S_T - K)^+}{K^2} dK + \int_0^{F_T} \frac{(K - S_T)^+}{K^2} dK .$$

It is clearly not convenient to trade in an infinite number of contracts, hence the above expression needs to be approximated by using finitely many European options. Assume we are able to trade in options with strikes $0 < K_{-n_p} < \dots < K_0 < \dots < K_{n_c}$ where $K_0 = F_T$ is ATM. Since $\log x$ is concave, we can construct a following sub-hedge on the interval $[K_{-n_p}, K_{n_c}]$:

$$\log S_T \geq \log F_T + \sum_{k=1}^{n_c} w_k^c (S_T - K_{k-1})^+ + \sum_{k=1}^{n_p} w_k^p (K_{-(k-1)} - S_T)^+ ,$$

with appropriate weights w^c and w^p , cf. appendix B. A super-hedge can be constructed in a similar way.

A common point of concern is the fact that this super-hedge needs to be cut off at some point. The above formula is very sensitive to price of OTM options, which in turn are usually very illiquid. Hence, the price of a variance swap can not usually be read from the price of traded options, even under the idealizing assumptions of deterministic rates, no credit risk and no cash dividends.

In practise, of course, there are many more effects which are not taken into account by such a simplifying framework, chiefly among them riskiness of interest rates.

Interpretation: Playing Implied vs. Realized Volatility

Consider the following idea: let us assume that we want to hedge an European option with payoff $H \geq 0$ using the Black&Scholes model: we fix some “set volatility” K and compute, at any time t , the price of the option as $\Pi_t^{\text{BS}}(S_t)$ using the volatility K . We then short the option and delta-hedge according to Black&Scholes’s delta $\Delta_t^{\text{BS}}(S) := \partial_S \Pi_t^{\text{BS}}(S)$, again computed with our “set volatility” K .

Assuming the real stock price is a diffusion

$$\frac{dS_t}{S_t} = (r_t - \mu_t) dt + \sigma_t dW_t$$

with true (stochastic) volatility σ , Ito shows that

$$d\Pi_t^{\text{BS}}(S_t) = \Delta\Pi_t^{\text{BS}} dS_t + \frac{1}{2}\Gamma_t^{\text{BS}} \sigma_t^2 S_t^2 dt + \theta_t^{\text{BS}} dt$$

where we used Gamma $\Gamma_t^{\text{BS}}(S) := \partial_{SS}^2 \Pi_t^{\text{BS}}(S)$ and Theta $\theta_t^{\text{BS}}(S) := \partial_t \Pi_t^{\text{BS}}(S)$. Since Π satisfies the Black&Scholes-PDE, we have that

$$\theta_t^{\text{BS}}(S) = -\frac{1}{2}\Gamma_t^{\text{BS}} K^2 S^2 ,$$

hence we get the standard result that our accumulative hedging error is

$$\begin{aligned} \text{P\&L}_{\text{BS}}(H) &:= -\left(H(S_T) - \Pi_0(S_0)\right) + \int_0^T \Delta_t^{\text{BS}} dS_t \\ &= \frac{1}{2} \int_0^T \Gamma_t^{\text{BS}} S_t^2 (\sigma_t^2 - K^2) dt . \end{aligned}$$

That means that under the assumption that the real stock price is a diffusion, we can bet on a discrepancy between implied squared volatility K^2 and true squared volatility σ^2 : if we assume, for example, that the quoted implied volatility of some call is probably going to be below the true volatility of the asset, we could buy that call and delta-hedge it with its initial implied volatility. Since the call's gamma is positive, $\text{P\&L}_{\text{BS}}$ will be positive if the true volatility is indeed always higher than the implied volatility.

This approach, however, has one drawback: Cash Gamma $\Gamma_t^{\text{BS},\$} := \Gamma_t^{\text{BS}} S_t^2$ of a call is not a constant. Hence, the gain we make if we are right in our prediction will depend a lot on *when* the true volatility is above the implied volatility. This is particularly a concern if the spot moves away from the strike of the call in which case gamma decays rapidly and the effect of our bet diminishes. It also means that even if we are on average right that the true volatility is above the implied volatility, we may well still make a loss with the above strategy.

To alleviate this drawback, we would therefore require an option whose Cash Gamma is independent of the spot level, for example a constant. This payoff now is exactly the log-contract, $H(S_T) := -2 \log S_T$, i.e.

$$\text{P\&L}_{\text{BS}}(2 \log) = \int_0^T (\sigma_t^2 - K^2) dt .$$

Moreover, the delta of a log-contract *in Black-Scholes* is easily seen to be exactly $\Delta_t^{\text{BS}} = P(t, T)/S_t$ regardless of the implied volatility. This is in line with our formula (28).¹² In this sense, a variance swap can be seen as the delta-hedged European option position which provides a constant exposure in cash to the discrepancy between implied and true volatility.

REMARK 3.4 (Delta of a Variance Swap in Black-Scholes) *In Black-Scholes, a variance swap which pays pure quadratic variation has no delta – its payoff*

¹²More details on this construction can be found in [5].

is a constant. However, the delta of a short position in such a variance swap, plus a long position in its log-hedge and the appropriate funding in cash (see below) will be $2P(T, t)/S_t$.

However, even in Black-Scholes a standard variance swap paying realized variance will have a residual daily delta which arises from the daily log-squared returns which comprise its payoff, as can be seen in (30).

REMARK 3.5 (Delta of a Variance Swap in Stochastic Volatility models) *In other volatility models, such as Dupire's implied local volatility, a variance swap will have some delta arising from the additional dependence of volatility on spot. In stochastic volatility models such as Heston's, a variance swap has only a very small delta which arises from the fact that it pays the realized variance rather than quadratic variation. If it were to pay just the latter, it wouldn't have a delta in the classical sense, but a "vega" towards a volatility which is highly correlated with spot.*

3.2 Cash Dividends

The presence of cash dividends alters the argumentation above only slightly. The main difference is that then (24) on page 16 does not hold since it omits the jumps due to the dividends. Assuming that X itself has no jumps, we obtain

$$\begin{aligned} \log \frac{S_{t_i}}{S_{t_{i-1}}} &= \int_{t_{i-1}}^{t_i} \frac{dS_t}{S_t} - \frac{1}{2} \int_{t_{i-1}}^{t_i} \frac{d\langle S \rangle_t}{S_t^2} \\ &\quad + \log \frac{S_{t_i}}{S_{t_i-}} - \left(\frac{S_{t_i}}{S_{t_i-}} - 1 \right) - \left(\frac{S_{t_i}}{S_{t_i-}} - 1 \right)^2, \end{aligned}$$

which has also been discussed in [10] and [2]. Following (1) on page 2 we can write at each dividend date

$$S_{\tau_j-} = \frac{S_{\tau_j} + \alpha_j}{1 - \beta_j}. \quad (33)$$

(we here use the assumption that dividends are paid right before the open). Let us also introduce the function

$$D_j(s) := \log \frac{s(1 - \beta_j)}{s + \alpha_j} + \frac{s\beta_j + \alpha_j}{s + \alpha_j} - \left(\frac{s\beta_j + \alpha_j}{s + \alpha_j} \right)^2.$$

Then, we can summarize just as above that the payoff of a variance swap which is not adjusted for dividends is given as

$$\text{RV}^I(T) = -2 \log \frac{S_T}{S_0} + 2 \sum_{i=1}^n \left(\frac{S_{t_i}}{S_{t_{i-1}}} - 1 \right) + 2 \sum_{j: \tau_j \leq T} D_j(S_{\tau_j}). \quad (34)$$

The key in this expression is that the additional term on the right side is simply a strip of European payoffs, which in turn can again be approximated by (32).

If dividends are taken out of the computation of the variance swap returns, the formula changes accordingly to

$$RV^S(T) = -2 \log \frac{S_T}{S_0} + 2 \sum_{i=1}^n \left(\frac{S_{t_i}}{S_{t_{i-1}}} - 1 \right) + 2 \sum_{j: \tau_j \leq T} D_j^S(S_{t_i}) ,$$

where

$$D_j^S(s) := \log \frac{s(1 - \beta_j)}{s + \alpha_j} + \frac{s\beta_j + \alpha_j}{s + \alpha_j} ,$$

i.e. we omit the square term from D .

REMARK 3.6 (Are variance swaps cheaper with cash dividends?)

A very common question will be to ask whether variance swaps become more or less expensive if cash dividends are introduced.

It needs to be stressed that one can not simply price a variance with and without cash dividends off the same implied volatility surface: in the light of our discussion above it should be clear that changing the dividend assumption also changes the implied vol surface – or, in other words, if the implied vol surface is marked with one dividend assumption, then we should not simply price a variance swap with another assumption.

By experience (and re-fitting of the vol surface), however, the prices of variance swaps will fall if cash dividends are introduced. CHECK THIS WITH TRADING

3.3 A Note on the Incorporation of Credit Risk

The proper incorporation of credit risk into variance swap pricing in our framework is actually quite straight forward: by definition of the contract, realized variance will become infinite, hence the contract's payoff will simply be the maximum capped amount, $150\%K^2$, cf. remark 3.1 on page 16.

In other words, following a suggestion in [2], we can write the single stock variance swap payoff as

$$1_{\tau > T} \left(\frac{252}{n} RV^S(T) - K^2 \right) + 1_{\tau < T} 150\%K^2 ,$$

which can be achieved easily using the methods described above.

However, in practise one should keep in mind the typically very high correlation between volatility and default intensities. Hence, if modelling variance swaps in single stocks, it is advisable to incorporate stochastic default rates. The above will also not work if variance spikes after the inception of the product. In this case, the cap will have true optionality and needs to be evaluated with elaborate methods, see [10].

3.4 Related Products

The pricing of variance swaps can be extended to a series of products with very similar characteristics such as “Gamma swaps” and “Conditional” or “Corridor

Variance Swaps". The idea in all cases is the same, as the payoff is now

$$\frac{1}{n} \sum_{i=1}^n \left\{ 252 f(S_{t_{i-1}}) \left(\log \frac{S_{t_i}}{S_{t_{i-1}}} \right)^2 - g(S_{t_{i-1}}) K^2 \right\} . \quad (35)$$

For example, a variance swap has $f(s) = 1$ and $g(s) = 1$, while a *gamma swap* has $f(s) := s$ and $g(s) = 1$. A *conditional variance swap* with a lower barrier of B has $f(s) = 1_{s>B}$ and $g(s) = 1_{s>B}/n$, i.e. it is a variance swap which only pays out realized variance vs. set variance for those days where the stock price trades above the barrier. A *corridor variance swap*, finally, is given by $f(s) = 1_{s>B}$ and $g(s) = 1$, i.e. the fixed leg is always paid while the floating leg is only paid if the stock price trades above the barrier B . See Overhaus et al. [10] for some more information.

From a risk-management point of view, this product again separates into two: the right hand side is a strip of European options, mostly digitals, for which most desks have a well-developed pricing methodology.

The left hand part of the payoff can be evaluated as follows: first of all, let us ignore the coefficient $252/n$, i.e. we consider only

$$\text{RV}_f(T) := \sum_{i=1}^n f(S_{t_{i-1}}) \left(\log \frac{S_{t_i}}{S_{t_{i-1}}} \right)^2 .$$

As before, assume absence of any credit risk and that X has no jumps. Also, chose a function F such that

$$F''(s) = \frac{f(s)}{s^2} .$$

We can then write

$$\begin{aligned} \mathbb{E}[\text{RV}_f(T)] &= \mathbb{E} \left[\int_0^T \frac{f(S_{t-})}{S_{t-}^2} d\langle S \rangle_t \right] \\ &= 2\mathbb{E}[F(S_T) - F(S_0)] - 2\mathbb{E} \left[\int_0^T F'(S_{t-}) dS_t \right] - 2 \sum_{j: \tau_j \leq T} D_j^f(S_{\tau_j}) \end{aligned}$$

with a “dividend adjustment” term due to (33),

$$\begin{aligned} D_j^f(s) &:= \left\{ F(s) - F\left(\frac{s + \alpha_j}{1 - \beta_j}\right) \right\} \\ &\quad - F'\left(\frac{s + \alpha_j}{1 - \beta_j}\right) \left(\frac{s\beta_j + \alpha_j}{1 - \beta_j} \right) \\ &\quad - \frac{1}{2} f\left(\frac{s + \alpha_j}{1 - \beta_j}\right) \left(\frac{s\beta_j + \alpha_j}{s + \alpha_j} \right)^2 . \end{aligned}$$

For the example of a corridor (or conditional) variance swap, we have $f(s) := 1_{s>B}$, i.e. $F''(s) := 1_{s>B}/s^2$, which gives

$$F'(s) = \left(\frac{1}{B} - \frac{1}{s} \right)^+ \quad \text{and} \quad F(s) = \frac{1}{B} (s - B)^+ - \left(\log \frac{s}{B} \right) 1_{s>B} .$$

In fact, these terms are also very intuitive: they essentially postulate that no delta-hedging or log-contract type European option position is required below the barrier B . Above the barrier, the hedge and the European position are essentially the (shifted) quantities of the variance swap hedge.

To evaluate $\mathbb{E}[F(S_T) - F(S_0)]$, we can employ approximation via Europeans as before. The delta-hedging term can be approximated as follows:

$$\begin{aligned} \mathbb{E} \left[\int_0^T F'(S_{t-}) dS_t \right] &\approx \mathbb{E} \left[\sum_{i=1}^n (F(S_{t_{i-1}}) S_{t_{i-1}}) \left(\frac{S_{t_i}}{S_{t_{i-1}}} - 1 \right) \right] \\ &\approx \sum_{i=1}^n \mathbb{E} [F(S_{t_{i-1}}) S_{t_{i-1}}] \left(\frac{F_{t_i}}{F_{t_{i-1}}} - 1 \right), \end{aligned}$$

i.e. it is again a matter of evaluating a string of European options. In the case of a corridor variance swap, we have to evaluate a sum of call options struck at the barrier,

$$\sum_{i=1}^n \frac{1}{B} \mathbb{E} \left[\left(S_{t_{i-1}} - B \right)^+ \right] \left(\frac{F_{t_i}}{F_{t_{i-1}}} - 1 \right).$$

Finally, the sum of the dividend terms $D_j^f(S_{\tau_j})$ is once again a simple strip of European options.

The other payoff variants are equally straight forward.

REMARK 3.7 *Since all payoffs of the type (35) can essentially be priced with European options, their price coincides with the price computed using an implied local volatility framework.*

4 Summary

We have discussed a simple and theoretically sound framework for handling cash dividends and credit risk in equity modelling. We have shown the impact of such an approach on the pricing of European options, on implied volatility, how to derive the appropriate PDEs and we have shown how to incorporate dividend effects in the pricing and risk-management of variance swaps and related products.

While our framework makes no claim of sufficiency for today's very complex and highly developed markets, it is clearly a step forward towards a sound integration of cash dividends into any risk management system. Its numerical closeness to the standard "drift-only" methods in equity pricing make it particularly appealing.

A Ito for jump processes

Let X and Y be two random cadlag processes with finite activity jumps (for example, a diffusion mixed with a Poisson-process). Denote by X^c and Y^c the

continuous versions of X and Y , respectively, i.e. $dX^c = dX_{t-}$. Then,

$$\begin{aligned}
dF(t, X_t, Y_t) &= \partial_t F(t, X_{t-}, Y_{t-}) dt \\
&\quad + \partial_x F(t, X_{t-}, Y_{t-}) dX_t^c + \partial_y F(t, X_{t-}, Y_{t-}) dY_t^c \\
&\quad + \frac{1}{2} \partial_{xx} F(t, X_{t-}, Y_{t-}) d\langle X^c \rangle_t + \frac{1}{2} \partial_{yy} F(t, X_{t-}, Y_{t-}) d\langle Y^c \rangle_t \\
&\quad + \partial_{xy} F(t, X_{t-}, Y_{t-}) d\langle Y^c, X^c \rangle_t \\
&\quad + \sum_t \left\{ F(t, X_t, Y_t) - F(t, X_{t-}, Y_{t-}) \right\}
\end{aligned}$$

If X and Y have jumps of infinite activity, then the above formula does not hold and the following reformulation has to be used

$$\begin{aligned}
dF(t, X_t, Y_t) &= \partial_t F(t, X_{t-}, Y_{t-}) dt \\
&\quad + \partial_x F(t, X_{t-}, Y_{t-}) dX_t + \partial_y F(t, X_{t-}, Y_{t-}) dY_t \\
&\quad + \frac{1}{2} \partial_{xx} F(t, X_{t-}, Y_{t-}) d\langle X^c \rangle_t + \frac{1}{2} \partial_{yy} F(t, X_{t-}, Y_{t-}) d\langle Y^c \rangle_t \\
&\quad + \partial_{xy} F(t, X_{t-}, Y_{t-}) d\langle Y^c, X^c \rangle_t \\
&\quad + \sum_t \left\{ F(t, X_t, Y_t) - F(t, X_{t-}, Y_{t-}) \right\} \\
&\quad - \sum_t \left\{ \partial_x F(t, X_{t-}, Y_{t-}) (X_t - X_{t-}) + \partial_y F(t, X_{t-}, Y_{t-}) (Y_t - Y_{t-}) \right\}
\end{aligned}$$

B Replication of Convex Payoffs with European Options

For any convex function f , let

$$h_1(x) := f(x_0) + w_1(x - x_0), \quad \tilde{w}_1 := \frac{f(x_1) - f(x_0)}{x_1 - x_0}$$

such that $h_1 \geq f$ on $[x_0, x_1]$ and $h_1(x_1) = f(x_1)$. On $[x_1, x_2]$, set

$$h_2(x) := f(x_1) + \tilde{w}_2(x - x_1) = f(x_0) + \tilde{w}_1(x - x_0) + (\tilde{w}_2 - \tilde{w}_1)(x - x_1), \quad \tilde{w}_2 := \frac{f(x_2) - f(x_1)}{x_2 - x_1}$$

Iteration and using

$$w_k := \frac{f(x_k) - f(x_{k-1})}{x_k - x_{k-1}} - w_{k-1}$$

with $w_0 := 0$ yields that

$$h(x) := f(x_0) + \sum_{k=1}^n w_k(x - x_{k-1})^+$$

dominates f on $[x_0, x_n]$.

On the downside, we apply the same trick:

$$g(x) := f(x_0) + \tilde{w}'_1(x - x_0) = f(x_0) - \tilde{w}'_1(x_0 - x) , \quad \tilde{w}'_1 := \frac{f(y_1) - f(x_0)}{y_1 - x_0}$$

such that

$$h(x) := f(x_0) - \sum_{k=1}^m w'_k(x_{k-1} - x)^+$$

with

$$w'_k := \frac{f(y_k) - f(y_{k-1})}{y_k - y_{k-1}} - w'_{k-1}$$

References

- [1] *H.Buehler*:
“Volatility Markets: Consistent modeling, hedging and practical implementation”, PhD thesis TU Berlin, 2006
<http://www.quantitative-research.de/dl/HansBuehlerDiss.pdf>
- [2] *H.Buehler*:
“Options On Variance: Pricing And Hedging” (presentation), IQPC Volatility Trading Conference, London, 2006
<http://www.quantitative-research.de/dl/IQPC2006-2.pdf>
- [3] *H.Buehler*:
“Expensive Martingales”, Computational Finance, 2006
- [4] *F.Delbaen, W.Schachermayer*:
“A General Version of the Fundamental Theorem of Asset Pricing”, Mathematische Annalen 300, pp. 463-520, 1994
- [5] *K.Demetri, E.Derman, M.Kamal, J.Zou*:
“More Than You Ever Wanted To Know About Volatility Swaps”, Quantitative Strategies Research Notes, Goldman Sachs, 1999
- [6] *B.Dupire*:
“Pricing with a Smile”, Risk, 7 (1), pp. 18-20, 1996
- [7] *J.Gatheral*:
“The Volatility Surface: A Practitioner’s Guide”, Wiley 2006
- [8] *R.Lee*:
“The Moment Formula For Implied Volatility At Extreme Strikes”, pp. 469-480, Mathematical Finance, 14, 2004
- [9] *A.Neuberger*:
“Volatility Trading”, London Business School WP, 1992

- [10] *Overhaus et al.*:
“Equity Hybrid Derivatives”, Wiley, 2006