Optimal Weak Static Hedging of Equity & Credit Risk Using Derivatives

Dirk Becherer
Institut of Mathematics
Humboldt Universität zu Berlin
Unter den Linden 6
D-10099 Berlin, Germany
becherer(at)math.hu-berlin.de

Ian Ward*
Credit Portfolio Management, KPMG
1 Canada Square
London, E14 5AG, U.K.
ian.ward(at)KPMG.co.uk
& ian.ward(at)lambdat.com

revised version †of March 26, 2009

Abstract

We develop a generic method for constructing a weak static minimum variance hedge for a wide range of derivatives that may involve optimal exercise features or contingent cash flow streams, to provide a hedge along a sequence of future hedging dates. The optimal hedge is constructed using a portfolio of pre-selected hedge instruments which could be derivatives with different maturities. The hedge portfolio is weakly static in that it is initiated at time zero, does not involve intermediate re-balancing, but hedges may be gradually unwound over time. We study the static hedging of a convertible bond to demonstrate the method by an example that involves equity and credit risk. We investigate the robustness of the hedge performance with respect to parameter and model risk by numerical experiments.

Keywords: static hedging, minimum variance hedging, displaced diffusion, stochastic volatility, calibration, convertible bond

*Corresponding author

†We thank participants of the London Graduate School in Mathematical Finance Workshop, the IMA Conference on Computational Finance and the AMAMEF conference in Bedlewo, for references and suggestions, in particular T. Bielecki, R. Frey, A. Giese and J. Maruhn. As usual, we are responsible for any errors and the statements in the paper only reflect the author's personal viewpoints. Their professional affiliation does not change this.

1 Introduction

Consider an agent (the *hedger*) who holds a primary derivatives position with exposure to market risk factors like equity and credit risk. The derivatives may involve various optimal exercise and other exotic features, giving rise to risky contingent payoffs at a sequence of future dates. In hoping to achieve a hedge that is more robust with respect to model risk, the agent decides on a static hedging strategy and selects appropriate hedging instruments to capture perceived risk factors. Those instruments will typically include other derivatives of various type. Not expecting to find a perfect hedge, the agent seeks an optimal approximating one that makes the residual risk small. He must then compute the notionals for his hedges, covering all hedging dates, and assess the robustness of his hedge. This motivates the problems to be addressed in this paper.

In general, hedging a primary contingent claim means to replicate its payoff (approximatively) by trading other securities, in order to eliminate (reduce) the risk. The general property of *static* hedging is that it does not involve re-balancing, in contrast to dynamic hedging. In this paper, a portfolio is called static if its hedging positions are not re-balanced after portfolio initiation, but portfolio constituents may be sold over time, before they expire. hedging. A classical result on static hedging is that any European claim, whose payoff is a function of a (single) underlying, may be statically hedged to arbitrary accuracy using calls and puts of maturity matching the claim if all strikes are tradable, see Breeden and Litzenberger (1978). This (strong) static hedge is model risk free. Let us emphasise, that the static hedging methods discussed in the sequel are not model risk free, but all rely to a greater or lesser extent on model assumptions. Many articles have studied the static hedging of barrier options, see Andersen et al. (2002); Brown et al. (2001); Carr et al. (1998); Derman et al. (1995); Giese and Maruhn (2007); Nalholm and Poulsen (2006). Most results are concerned with the replication of the barrier option by static hedging with European vanilla options. To achieve this, they exploit the specific structure of the option and certain model assumptions like a deterministic future volatility surface, continuously evolving prices, or the availability of infinitely many vanilla options, cf. Joshi (2003). Hence, the results do not transfer easily to derivatives of other types, and incorporate model risk, see Nalholm and Poulsen (2006). For stochastic volatility models, the said assumptions are typically not satisfied and replication by vanilla options is no longer possible. This calls for another more generally applicable notion of hedging. For results on over-hedging (super-replication) under less restrictive model requirements, we refer to Brown et al. (2001); Andersen et al. (2002); Giese and Maruhn (2007), while references for approximative (minimum variance) hedging are given below. Less seems to have been written about static hedging for (and with) derivatives other than barrier and European options. The only result on static hedging of options involving optimal exercise features that we are aware of is in Jäckel and Rebonato (2001). Their model assumptions are markedly different from ours, assuming a deterministic volatility surface restricted to a single underlying.

In comparison to the cited literature, our aim here is to develop a *generic* static hedging method that is applicable to a wider range of derivatives, both as hedge instruments and also as primary derivatives to be hedged, for a broad range of models, without imposing particular assumptions. To achieve this we adopt for our purpose of weak static hedging the approximate hedging approach

by Bouleau and Lamberton (1989); Föllmer and Sondermann (1986), as a computationally feasible method that is quite generally applicable. This method can be seen as 'market-based' minimum variance hedging where the residual risk of the hedged portfolio is minimised under the (empirical) pricing measure; see Remark 2.1 for further comments. The examples in the above, in Föllmer and Schweizer (1990) or Lamberton and Lapeyre (1993) are on dynamic hedging of European options where the hedging instruments could be other European options of the same maturity. In contrast, our contribution is to develop a (weak) static version of minimum variance hedging with possibly more exotic derivatives for a sequence of multiple future hedge dates. To investigate the performance and the robustness of the static hedges with respect to parameter and model risk, we conduct what may be described as a controlled experiment. We assume knowledge of a hypothetical *real world* model against which the performance of the hedge can be measured, while the computation of the static hedge is performed in a different *hedger's* model. The robustness of the static hedge against parameter and model risk is demonstrated under variation of the hedge model parameters, and by evaluating the hedge performance in the *real world* model. Two qualitatively different but computationally tractable models are developed for the *hedge* and *real world* models. The hedger's model is an extension of the classical displaced diffusion model from Rubinstein (1983), that includes credit risk into a one factor model for both equity and credit risk, yet maintains closed analytic formulas for the risk neutral transition probabilities over distinct periods. The model is designed to have just enough parameters such that it can be calibrated to the term structures of given market prices of liquid, vanilla equity and credit derivative securities. In contrast, the model associated with the real world is a time homogeneous two factor model that can be viewed as an affine extension of Heston's stochastic volatility model to include credit risk. The qualitative difference puts the hedger's model at a disadvantage. This is intended to study the model robustness of the hedges. By numerical experiments we study the performance of the static hedge for a convertible bond. We show that this contingent claim may be robustly hedged by vanilla equity and credit derivative securities. The statically hedged position exhibits little parameter and model risk.

The paper is organised as follows. Section 2 describes a method for computing an optimal weak static hedge for a portfolio of primary instruments, by using backward induction over multiple hedge dates in combination with co-initial regression. Section 3 introduces the (DD+) model that will be used to compute a hedge. Section 4 describes the extended Heston (EH) model, representing the real world, and the pricing of instruments. Section 5 applies the method to hedge a convertible bond and assesses the hedge performance. Section 6 summarises and concludes, highlighting, in particular, that choosing appropriate hedge instruments is essential to achieve a robust weak static hedge.

2 The Weak Static Minimum Variance Hedge

2.1 The General Framework

By 'static hedging' we mean, throughout this paper, a portfolio of hedging instruments, including derivatives, that is initiated at time zero, does not involve

re-balancing and may only change by selling its constituent instruments, possibly before their expiry. In this sense, the hedging method can be called *weak static*, following the classification in Joshi (2003), chapter 10.7.

Below, we list general requirements of the method for the model, for the derivatives that can be hedged, and for the possible hedging instruments. Given assumptions A1-A4 are met, it is possible, in principle, to use the weak static hedge methods that are presented in the sequel.

A1. The information flow (in continuous time) is modelled by the filtration $(\mathcal{F}_t)_{t\geq 0}$ of a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$. With some finite time horizon T, the model has occurrence times $0 = t_0 < \cdots < t_K = T$ to include (in approximation) all derivative cash flow payment dates, hedging dates and maturities of calibration securities to be considered.

A2. The model is constructed from some state process $(\mathfrak{X}_t)_{t\in[0,T]}$ taking values in a polish space like $\mathbb{R}^d$, which is a discrete time Markov process at times t_k, i.e. $(\mathfrak{X}_{t_k})_{k=0,\ldots,K}$ is Markov with respect to $(\mathcal{F}_{t_k})_{k=0,\ldots,K}$.

A3. We consider (derivative) securities of some maturity t_n, whose structure permits computation of their no-arbitrage value at any time t_k as a conditional expectation under the (risk neutral) *pricing measure* $\mathbb{P}$ as follows

$$V_{t_k} = \begin{cases} f_k(\mathfrak{X}_{t_k}, \mathbb{E}[\hat{f}_k(\mathfrak{X}_{t_k}, V_{t_{k+1}})|\mathcal{F}_{t_k}]), & 0 \leq k < n, \\ f_n(\mathfrak{X}_{t_n}), & k = n, \end{cases} \tag{2.1}$$

for derivative specific functions f_k and $\hat{f}_k$.

A4. All involved securities prices V_{t_k} are square integrable (under $\mathbb{P}$).

By A2-3, $V_{t_k} = v_k(\mathfrak{X}_{t_k})$ for suitable functions v_k. Intuitively, (2.1) requires that the time-t_k value of a derivative's future cash flows depends on values of $\mathfrak{X}_t$ on or after t_k only. For most derivatives, like subsequent examples in this paper, $\hat{f}_k(\mathfrak{X}_{t_k}, V_{t_{k+1}}) = V_{t_{k+1}}$ is sufficient, but (2.1) is more general.

For example let $\mathfrak{X}_{t_k}$ include the stock price process S_{t_k} underlying a call option with strike K and expiry t_n. Then, the function from A3 is

$$f_k = \begin{cases} (S_{t_k} - K)^+, & : \ k = n \\ \max\left(e^{-r_k(\mathfrak{X}_{t_k})(t_{k+1}-t_k)}\mathbb{E}[V_{t_{k+1}}|\mathcal{F}_{t_k}],\, 1_A(S_{t_k} - K)^+\right) & : \ k < n. \end{cases} \tag{2.2}$$

where $1_A = 1$ if the option is American (Bermudan type) and 0 if European, with r_k denoting the interest rate. Barrier options fit into the framework when the barrier is discretely monitored (in approximation). Similarly, one can deal with a unit currency convertible bond of maturity t_n, fixed coupon c paid at dates $Coup = \{t_{c_i},\, i = 1, \ldots, N_c\}$, accruing over time Δt, that may be converted to ζ stocks on dates $Conv = \{t_{v_j},\, j = 1, \ldots, N_v\}$, using

$$f_k = \begin{cases} \max(1 + c\Delta t, \zeta S_{t_n}) 1_{\{\tau > t_n\}} + R 1_{\{t_{n-1} < \tau \leq t_n\}} & : \ k = n, \\ \max\Big(e^{-r_k(t_{k+1}-t_k)}\mathbb{E}[V_{t_{k+1}}|\mathcal{F}_{t_k}] + c\Delta t 1_{\{t_k \in Coup\}}, & \\ \qquad \zeta S_{t_k} 1_{\{t_k \in Conv\}}\Big) 1_{\{\tau > t_k\}} + R 1_{\{t_{k-1} < \tau \leq t_k\}} & : \ k < n, \end{cases} \tag{2.3}$$

where τ denotes the time of default of the bond issuer, with suitable $\mathfrak{X}_t = (S_t, 1_{t<\tau})$. Now consider a credit default swap of unit currency notional where protection premium coupon s is paid on dates t_{c_i}, accruing over time Δt, for protection against default losses $(1-R)$. Under the simplifying assumption that default payments are made at the next model time after default, one can take

$$f_k = \begin{cases} -s\Delta t 1_{\{\tau>t_n\}} + (1-R)1_{\{t_{n-1}<\tau\le t_n\}} & : \ k=n, \\ \left(e^{-r_k(t_{k+1}-t_k)}\mathbb{E}[V_{t_{k+1}}|\mathcal{F}_{t_k}] - s\Delta t 1_{\{t_k\in Coup\}}\right) 1_{\{\tau>t_k\}} & \\ \qquad +(1-R)1_{\{t_{k-1}<\tau\le t_k\}} & : \ k<n. \end{cases} \tag{2.4}$$

Forward starting and cliquet options are derivatives with moderate path dependence in that they are dependent on the history of the underlying at only a small number of times. To see that they fit with A3, consider forward starting options with maturities t_n whose strike are struck ATM at time $t_{n-\ell}$, with an offset of $\ell \ge 1$ periods ($\ell \le n$). With Markovian stock price (S_{t_k}) and $\mathfrak{X}_{t_k} = (S_{t_{(k-\ell)\wedge 0}}, \ldots, S_{t_k})$, the forward starting call paying $(S_{t_n} - S_{t_{n-\ell}})^+$ at maturity t_n is valued consistently with A3 by $f_k = e^{-r_k(t_{k+1}-t_k)}\mathbb{E}[V_{t_{k+1}}|\mathcal{F}_{t_k}]$ at t_k for $k < n$. Derivatives with more path dependence such as Asian options whose payoff is based on some averaging process A_t of an underlying process S_t, may satisfy specifications A2,3 by enlarging the state space, e.g. to $\mathfrak{X}_t = (S_t, A_t)$.

In the sequel we consider a portfolio of one or more primary instruments that is to be statically hedged by a suitable portfolio of preselected instruments. It is natural, though not necessary, to think of the hedge instruments as liquid vanilla derivatives which are less exotic than the primary instruments to be hedged. In later sections, hedge instruments will be vanilla derivatives that are European call and put options, credit default swaps and zero coupon bonds. One could also use other hedge instruments, e.g. American or forward starting options. The method to be suggested does not require that the hedging instruments expire at the hedging date for which they are employed, and allows different forward start dates of the hedge instruments. Yet, the examples in the paper will be predominantly spot starting options which expire at the hedging dates to which they are associated.

Remark 2.1. Since our task is *hedging* exclusively our model is specified under the pricing measure $\mathbb{P}$, understanding that $\mathbb{P}$ is to be inferred by calibration to present market prices (see examples later). Given that the real market has chosen the pricing measure for our model, the pricing problem is solved a priori, as noted in Bouleau and Lamberton (1989). To find on optimal approximative hedge, we use the criterion from Föllmer and Sondermann (1986) and Bouleau and Lamberton (1989). Minimising the hedging P/L variance under $\mathbb{P}$ is a tractable criterion that permits hedge computations using a range of possibly complex derivatives. The residual hedging error that remains, being small under the pricing measure in a mean square sense, could be hedged dynamically afterwards. If one would start with $\mathbb{P}$ as the 'statistical' real probability measure, the valuations in (2.1) would involve the stochastic discount factor (pricing kernel). But for typical models, the real dynamics (drifts) of the processes would be hard, if not impossible, to estimate reliably.

2.2 The Weak Static Hedging Method

To fix notation, let us first describe the simple minimum variance hedge for a single horizon t_k, which will serve as a building block for the static hedge that uses a strip of derivatives to hedge a given exotic at multiple future hedge dates.

Let Π_k denote a portfolio comprising N_k primary instruments at time t_k. Let H_k be the time-t_k value of this portfolio and H_k^i, $i \leq N_k$, the value of each primary instrument. Assume that all the primary instruments are of maturity greater than or equal to t_k. By A2-3 the value of this portfolio may be expressed as a function of the Markov process $\mathfrak{X}$ at time t_k, i.e.

$$H_k(\mathfrak{X}_{t_k}) = \sum_{i=1}^{N_k} H_k^i(\mathfrak{X}_{t_k})\,. \tag{2.5}$$

At time t_k, let d_k hedge instruments be available of maturity t_k or greater, with values $Y_k^i = Y_k^i(\mathfrak{X}_{t_k})$ for $i \leq d_k$. Let $\theta_k^i > 0$ (or < 0) denote a long (or short) position in the ith hedge instrument. The hedging portfolio then has value

$$Y_k(\mathfrak{X}_{t_k}) = \sum_{i=1}^{d_k} \theta_k^i Y_k^i(\mathfrak{X}_{t_k}) = \theta_k.Y_k\,. \tag{2.6}$$

For illustration, take interest rates as deterministic with average rate of interest over time interval $[0,t]$ being $R(0,t)$. The objective of the static hedge is to select $\theta_k \in \mathbb{R}^{d_k}$ to minimise the variance of the total value (profit and loss P/L) $H_k + \theta_k.Y_k - e^{R(0,t_k)t_k}(H_0 + \theta_k.Y_0)$ of the hedged portfolio at time t_k, where the last term comprises the cost of financing, with $H_0 = e^{-R(0,t_k)t_k}\mathbb{E}[H_k]$ and $Y_0 = e^{-R(0,t_k)t_k}\mathbb{E}[Y_k]$. With $\Delta Y_k := Y_k - \mathbb{E}[Y_k]$ and $\Delta H_k := H_k - \mathbb{E}[H_k]$ one can write the time-t_k profit and loss compactly as $\Delta H_k + \theta_k.\Delta Y_k$. The hedge portfolio θ_k is chosen to solve the minimisation problem

$$\theta_k = \arg\min_{\theta\in\mathbb{R}^{d_k}} \mathbb{E}[(\Delta H_k + \theta.\Delta Y_k)^2] = \min_{\theta\in\mathbb{R}^{d_k}} \mathrm{Var}[(H_k + \theta.Y_k)]\,. \tag{2.7}$$

The first-order condition for the minimum yields a system of linear equations

$$\mathrm{Cov}(\Delta Y_k)\theta_k + \mathrm{Cov}(H_k, Y_k) = 0 \tag{2.8}$$

to be solved for θ_k. In (2.8), $\mathrm{Cov}(\Delta Y_k)$ is the $d_k \times d_k$-covariance matrix of ΔY_k, and $\mathrm{Cov}(\Delta H_k, \Delta Y_k)$ is the covariance vector $\left(\mathbb{E}[\Delta H_k \Delta Y_k^i]\right)_{i=1,\ldots,d_k}$. Assuming $\mathrm{Cov}(\Delta Y_k)$ is non-singular one may solve for unique optimal hedge notionals θ_k using standard methods, once the covariances in (2.8) are computed. If $\mathrm{Cov}(Y_k^{(k)})$ is singular then some t_k-hedging instruments are redundant, and one cannot expect the hedge to be unique. Note that the static hedging notional θ_k is linear with respect to the risky position H_k that is to be approximated.

We consider now the static minimum variance hedge of a portfolio Π of primary instruments at multiple future hedging horizons. With reference to the single hedge date results above, let us now state the general existence and uniqueness result for the static hedge portfolio of this work.

Proposition 2.2. Assume A1-4. For hedging horizons $0 < T_1 \cdots < T_M$, being a subset of the times from $A1$, there exist $\theta_k^* \in \mathbb{R}^{d_k}$, $1 \leq k \leq M$, that simultaneously solve the (weak) static minimum variance hedging problem

$$\min_{\theta_k \in \mathbb{R}^{d_k}} \mathrm{Var}\left[H_k + \sum_{j>k} \theta_j . Y_k^{(j)} + \theta_k . Y_k^{(k)} \right] \qquad \text{for all } k = 1 \dots M \tag{2.9}$$

for a given primary instrument with value process $(H_k)_{k=1,\dots,M}$ and $d_j \in \mathbb{N}$ preselected time-T_j hedge instruments, $j = 1, ..., M$, with value processes $(Y_k^{(j)})_{k=1,\dots,j}$ in $\mathbb{R}^{d_j}$. For any k, the minimising argument θ_k^* for (2.9) exists and $\theta_k^* Y_k^{(k)}$ is unique a.s.. The hedging notionals θ_k^* are unique if and only if the covariance matrix $\mathrm{Cov}(Y_k^{(k)})$ for the time-T_k hedge instruments is non-singular.

The variance minimising portfolio is found by solving systems of linear equations like (2.8) by backward induction for $T_M, \dots, T_1$, to construct the solution to the co-initial regression problems (2.9) for multiple horizons. For comparison, static variance minimisation for a single time horizon (Breeden and Litzenberger (1978)) uses simple linear regression but no backward induction, while dynamic minimum variance hedging (Föllmer and Schweizer (1990)) involves sequential regression. Though mathematically straightforward to obtain, the proposed weak static notion of minimum variance hedging for multiple horizons is new to the literature, to the best of our knowledge. We note that Derman et al. (1995), Jäckel and Rebonato (2001) and Joshi (2003) contain similar ideas, in that they also use backward induction to construct a static hedge. The main difference is, that they aim for (perfect) replication, while we only seek to minimise the variance of the hedging error to obtain an optimal approximate hedge. Such a weaker notion of hedging is useful since static replication is only possible under particular assumptions (cf. section 1), not in general. Since variances are model-dependent, robustness with respect to model risk is an issue to be considered in later sections.

Proof. With $h := H_M - \mathbb{E}[H_M]$ and $x := Y_M^{(M)} - \mathbb{E}[Y_M^{(M)}]$ the Mth minimisation problem is

$$\theta_M^* = \arg\min_{\theta \in \mathbb{R}^{d_M}} ||x^T\theta - h||_2^2 = \arg\min_{\theta \in \mathbb{R}^{d_M}} \mathbb{E}[(x^T\theta - h)^2]. \tag{2.10}$$

For $x \in L^2(\mathbb{R}^{d_k})$, the space $\mathcal{K}^k = \{x^T\theta : \theta \in \mathbb{R}^{d_k}\}$ is a finite dimensional (hence closed) subspace of the Hilbert space of square integrable random variables. So there is a unique orthogonal projection $x^T\theta_M^* \in \mathcal{K}^M$ of h onto $\mathcal{K}^M$, solving (2.10). The existence of the minimising argument for (2.9) follows by backward induction. If θ_i^* exist for $k < i \leq M$ then the kth minimisation problem for $\theta_k^* \in \mathbb{R}^{d_k}$ may be written in the form of (2.10) with $h = \widetilde{H}_k - \mathbb{E}[\widetilde{H}_k]$ and $x = Y_k^{(k)} - \mathbb{E}[Y_k^{(k)}]$ for known square integrable $\widetilde{H}_k = H_k + \sum_{i>k} \theta_i . Y_i^{(k)}$. So by induction θ_k^* exist for all $k = M, ..., 1$, minimising (2.9).

To prove the uniqueness condition, suppose θ^* and $\theta^* + \delta$ in $\mathcal{K}^k$ minimise (2.9) for some k. By uniqueness of the orthogonal projection of h onto $\mathcal{K}^k$, it holds that $x^T\theta^* = x^T(\theta^* + \delta)$, and so

$$0 = ||x^T\delta||_2^2 = \mathbb{E}[\delta^T(xx^T)\delta] = \delta^T \mathrm{Cov}(x)\delta\,.$$

In view of (2.8), it follows that $\delta \neq 0$ is only possible when $\mathrm{Cov}(x)$ is singular. □

The static hedging idea can be described as follows. Hold a suitable position in hedge instruments of maturity T_M or greater to minimise the variance of profit and loss from the primary instruments at time T_M. Perform a backward valuation to obtain the value of the primary and hedge instruments at time T_{M-1} and use *pre-selected* hedge instruments available for maturity T_{M-1} to minimise the profit and loss of this new portfolio. Repeat this procedure until the closest hedge date to today is reached. The initial value of the statically hedged portfolio Π given $\mathfrak{X}_{t_0} = x$ is given by $\Pi(x) = H_0(x) + \sum_{j=1}^{M} \theta_j . Y_0^{(j)}(x)$. The hedge portfolio is set up at time zero and the hedge instruments are gradually unwound at the hedging times that they are associated to, possibly before their expiry. Intuitively, hedge instruments for hedge date t_k are *pre-selected* to hedge contingent cash flows or exercise optionalities close to the time interval $[t_k, t_{k+1})$ where those originate. Clearly, hedge instruments need to capture major risk factors of the primary instrument; this should include the timing of contingent cash flows and exercise optionalities. European options, e.g., may serve as hedging instruments for primary instruments with a suitable weak form of paths dependence (e.g. barrier or American options with otherwise European payoff profiles); It would then be natural to use those European options with maturity t_k as hedge instruments for that hedge date.

To answer a potential question: for any hedge date, why not consider all available unexpired European options as hedge instruments? Then, theoretically, the variance of the hedge error (2.9) would be minimised over *ALL* hedge instruments. Such a minimisation problem is likely to become mathematically ill-posed. Practically, the number of hedge instruments should be reasonably small. And those hedges should be selected that are closest to the risks that have not already been hedged (at later dates).

Remark 2.3. Implementation and calibration issues.

1. One needs to evaluate the (conditional) expectations in (2.1) and (2.8) for all instruments in question. Unless instruments in question have closed form pricing formulas valuations need to be obtained numerically. To compute the risk neutral transition probability kernels $\mathbb{P}(\mathfrak{X}_{t_{k+1}} \in dx | \mathfrak{X}_{t_k} = x)$ and $\mathbb{P}(\mathfrak{X}_{t_k} \in dx)$, it is computationally efficient if $\mathfrak{X}_{t_k}$ is such that these kernels are given in closed form. Section 3 uses this for a lattice implementation.

2. The backward valuation may be computed by analytic or numerical methods, or by Monte Carlo simulation for high dimensional problems. Monte Carlo is particulary straightforward in the special case of static hedging of European options for a single horizon, see e.g. Pellizzari (2005).

3. The problem (2.8) can be ill conditioned. The covariance matrix and vector computed will involve certain numerical errors, and hedging instruments may be highly correlated. Some regularisation should be used, like truncated singular value decomposition or Tikhonov regularisation.

3 The Extended Displaced Diffusion Model

We describe stock price dynamics by an extension of the displaced diffusion model by Rubinstein (1983). Stock prices evolve according to a generalised displaced diffusion model with time varying coefficients, which is further extended to incorporate default of the stock issuer. A simple variant, called DD+, with

piecewise constant parameters is used later as the basis for a model that serves as the *hedger's* model in this paper. Like the classical model by Rubinstein (1983), our DD+ extension maintains analytic tractability in periods with constant coefficients. For jump diffusion models in general, see also Merton (1990).

General Model Framework

Consider a model for equity and credit derivative securities based on a single underlying state variable, the (defaultable) stock price S_t. Assume that (S_t) is an adapted process defined on the complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions. $\mathbb{P}$ is taken as the risk neutral measure under which security prices discounted by the risk free savings account $B_t = \exp(\int_0^t r(s)ds)$ are martingales. Interest rates $r(t) \geq 0$ are deterministic. To define (S_t), consider first an auxiliary process $\mathcal{X}_t$ evolving as

$$d\mathcal{X}_t = \big(r(t) + \lambda_t)\big)\mathcal{X}_t dt + \tilde{\sigma}(t, \mathcal{X}_t)dW_t, \quad t > 0\,, \tag{3.11}$$

where $\mathcal{X}_0 = S_0 = S(0)$ is the initial spot price of the stock, W is a standard $(\mathbb{P}, \mathbb{F})$-Brownian motion and $\lambda_t = \lambda(t, \omega)$ is the predictable $(\mathbb{P}, \mathbb{F})$-intensity of a doubly stochastic Poisson process[1] N_t, also adapted to $\mathbb{F}$. The time-inhomogeneous coefficient $\tilde{\sigma}$ is assumed affine in $\mathcal{X}$ with deterministic parameters $\sigma(t)$, $q(t)$ given by $\tilde{\sigma}(t, \mathcal{X}_t) = \sigma(t)\big(q(t)\mathcal{X}(t) + (1 - q(t))F(t)\big)$ for some deterministic function $F(t)$ (to be specified later (3.15)). As such the dynamics of $\mathcal{X}$ are similar to the displaced diffusion dynamics in Rubinstein (1983). If $q(t)$ is piecewise continuous and $q(t) < 1$ for some t then $\mathbb{P}(\inf_{t\in[0,T]} \mathcal{X}_t \leq 0) > 0$.

Let us define stopping times $\tau_N := \inf\{t > 0 : N_t = 1 > N_0 \equiv 0\}$ associated to the first jump of N and $\tau_D := \inf\{t \in [0, T] : \mathcal{X}_t \leq 0\}$ for the first hitting time of zero by $\mathcal{X}_t$, with the usual convention $\inf \emptyset = \infty$. Then, let

$$\tau := \tau_D \wedge \tau_N \equiv \min\{\tau_D, \tau_N\}, \tag{3.12}$$

be the time of default of the stock issuer and define the *defaultable* stock price process as $S_t = \mathcal{X}_t 1_{\{\tau > t\}}$. Note that S_t/B_t is a $(\mathbb{P}, \mathbb{F})$-local martingale.

Remark 3.1. Further salient features of the model are: (i) There are two sources of default: jump to default, which occurs at the totally inaccessible time τ_N, and diffusion to default which occurs at the predictable time τ_D.

(ii) Parameter $q(t)$ varies the weight between normal and lognormal stock returns. Like in Rubinstein (1983), it can be seen as a *skew* parameter since it primarily controls the model's implied volatility skews.

(iii) One obtains a stock price model as in Andersen and Buffum (2003), when taking $q(t) \equiv 1$ and an intensity of the form $\lambda_t = a(t)f_\lambda(S_t)$ for $0 \leq t < \tau$, where $a(t)$ is a deterministic function and

$$f_\lambda(S_t) = (S_t/S_0)^{-p} \tag{3.13}$$

for some $p \geq 0$; the intensity is then inversely related to the stock price and the volatility of short-term spreads equals $(-)p$ times the equity volatility.

(iv) If intensities are not state dependent, i.e. $f_\lambda \equiv 1$, and q_k do not depend on k, our model is an extension of the classical displaced diffusion model

[1] That means, conditional on knowing the information $\mathcal{H} := \sigma(\{\lambda_t : t \geq 0\})$, the process N_t is still a Poisson process with (known) intensity (λ_t) under information flow $(\mathcal{F}_t \vee \mathcal{H})_{t\geq 0}$.

with 'skew' from Rubinstein (1983), that includes default and excludes negative values of the underlying, while maintaining closed form solutions for prices of European vanilla options of any maturity. To see this, note that such parameterisation is a special case of the DD+ specification below. By (3.17), $\mathcal{X}_t$ in $S_t = \mathcal{X}_t 1_{\{\tau > t\}}$ is then a shifted lognormal random variable, independent of τ.

The Extended Displaced Diffusion Model DD+

We are going to use a simple version (i.e, piecewise constant parameters) of the displaced diffusion dynamics as a model capable of matching valuations of options and credit default swaps to their market prices. We now specify this version and refer to it as the extended displaced diffusion model DD+. Let the time dependent coefficients of (3.11) be piecewise constant so that $r(t) = r_k$, $q(t) = q_k$, $\sigma(t) = \sigma_k$, $a(t) = a_k$ for $t \in (t_{k-1}, t_k]$. The intensity is also piecewise constant and taken to be a function of the stock price at the beginning of each period, i.e. $\lambda_t = a_k f_\lambda(S_{t_{k-1}})$ (see (3.13)). The auxiliary process $\mathcal{X}$ satisfies

$$d\mathcal{X}_t = (r_k + a_k f_\lambda(s_{k-1}))\mathcal{X}_t dt + \sigma_k(q_k \mathcal{X}_t + (1-q_k)F_t^k)dW_t \tag{3.14}$$

for $t \in (t_{k-1}, t_k]$ over the kth period, with $\mathcal{X}_{t_{k-1}} =: s_{k-1}$ and

$$F_t^k := F_{t_{k-1}} \exp\left\{(r_k + a_k f_\lambda(s_{k-1}))(t - t_{k-1})\right\}, \tag{3.15}$$

$$F_{t_{k-1}} := S_0 \exp\left\{\sum\nolimits_{j=1}^{k-1}(r_j + a_j))(t_j - t_{j-1})\right\}. \tag{3.16}$$

By choosing F_t^k thus, applying Ito's Lemma to $\widetilde{\mathcal{X}}_t := q_k \mathcal{X}_t + (1-q_k)F_t^k$ shows that $\widetilde{\mathcal{X}}_t$ is lognormal (conditionally on $\mathcal{X}_{t_{k-1}} =: s_{k-1}$), and satisfies

$$d\widetilde{\mathcal{X}}_t = (r_k + a_k f_\lambda(s_{k-1}))\widetilde{\mathcal{X}}_t dt + \sigma_k q_k \widetilde{\mathcal{X}}_t dW_t\,, \quad t \in (t_{k-1}, t_k]\,. \tag{3.17}$$

Therefore, with $y_t^{(k)} := \sigma_k q_k (W_t - W_{t_{k-1}}) - \sigma_k^2 q_k^2 (t - t_{k-1})/2$, and using (3.15) it follows

$$\widetilde{\mathcal{X}}_t = \widetilde{\mathcal{X}}_{t_{k-1}} e^{(r_k + a_k f_\lambda(s_{k-1}))(t - t_{k-1}) + y_t^{(k)}} = \widetilde{\mathcal{X}}_{t_{k-1}} \frac{F_t^k}{F_{t_{k-1}}} e^{y_t^{(k)}}. \tag{3.18}$$

Rearranging terms, it follows for $t \in (t_{k-1}, t_k]$ that

$$\begin{aligned} y_t^{(k)} &= \log\Big(\frac{\widetilde{\mathcal{X}}_t}{F_t^k}\Big) - \log\Big(\frac{\widetilde{\mathcal{X}}_{t_{k-1}}}{F_{t_{k-1}}}\Big) \\ &= \log\Big(q_k \frac{\mathcal{X}_t}{F_t^k} + 1 - q_k\Big) - \log\Big(q_k \frac{\mathcal{X}_{t_{k-1}}}{F_{t_{k-1}}} + 1 - q_k\Big) \qquad (3.19) \\ &= \log\Big(q_k \frac{\mathcal{X}_t}{F_{t_{k-1}} e^{(r_k + a_k f_\lambda(s_{k-1}))(t - t_{k-1})}} + 1 - q_k\Big) - \log\Big(q_k \frac{\mathcal{X}_{t_{k-1}}}{F_{t_{k-1}}} + 1 - q_k\Big) \\ &=: \hat{Y}_k(\mathcal{X}_{t_k}; \mathcal{X}_{t_{k-1}}), \qquad (3.20) \end{aligned}$$

for $0 < q_k < 1$, similar results can be obtained for $q_k \geq 1$ and $q_k \leq 0$. This relation between $\mathcal{X}_t$ and the (conditional) Gaussian $y_t^{(k)}$ is used in the implementation of the (DD+) model described in the sequel.

The DD+ Model Implementation

The DD+ model facilitates numerical treatment, since transition probabilities are available in a simple analytical form over each period. Those will be derived next, see (3.22) below, and then applied for a recombining lattice implementation of the (discretised) stock price dynamics (3.14). The lattice should be seen as numerical approximation to compute prices and hedges, rather than a stochastic model in its own right. At each time slice for each lattice node of possible levels of the stock price, transition probabilities to nodes at the next time slice are to be approximated accordingly.

Recall the discretisation of time $0 = t_0 < \cdots < t_K = T$. For later purposes these times can be thought of as (moderately many) expiries of calibration and hedge instruments. At each time t_k, we discretise the stock price distribution to take values $S^{(j)}_{t_k} > 0$ for $j = -N^d_k$ to N^u_k, increasing in j, and to be zero in default. Let us note here that, in order to incorporate default time information into the underlying Markov state, several states associated to default will be used later. Let $\Delta S^{(j)}_{t_k} = S^{(j+1)}_{t_k} - S^{(j)}_{t_k}$ denote the spacing of discretised prices. The approximating transition probabilities between lattice nodes from $S_{t_{k-1}} = x > 0$ to $S^{(j)}_{t_k}$ are denoted by and computed as

$$p_{k-1,k}\left(S^{(j)}_{t_k} \,|\, x\right) := \mathbb{P}\left(x^-_{j,k} < S_{t_k} < x^+_{j,k} \,|\, S_{t_{k-1}} = x\right) \tag{3.21}$$

with $x^-_{j,k} := S^{(j)}_{t_k} - \frac{1}{2}\Delta S^{(j-1)}_{t_k} > 0$ and $x^+_{j,k} := S^{(j)}_{t_k} + \frac{1}{2}\Delta S^{(j)}_{t_k} > 0$. Omitting $k-1, k$ (j) sub(super)-scripts for ease of notation, the transition probabilities can be written as $\mathbb{P}\left(x^-_k < \mathcal{X}_{t_k} < x^+_k, I^{\mathcal{X}}_k > 0, \tau_N > t_k \,|\, \mathcal{X}_{t_{k-1}} = x, \tau_N > t_{k-1}\right)$ with $I^{\mathcal{X}}_k = \inf_{t_{k-1}<u<t_k} \mathcal{X}_u$. Since conditional on $\mathcal{X}_{t_{k-1}} = x$, point process N_t has constant intensity $\lambda_k(x)$ over $(t_{k-1}, t_k]$, the processes N_t and $\mathcal{X}_t$ are conditionally independent on $(t_{k-1}, t_k]$, implying

$$p(S_{t_k}|x) = e^{-(t_k - t_{k-1})\lambda_k(x)} \mathbb{P}\left(x^-_k < \mathcal{X}_{t_k} < x^+_k, I^{\mathcal{X}}_k > 0 | \mathcal{X}_{t_{k-1}} = x\right). \tag{3.22}$$

Transition probability (3.22) can be computed, by using (3.19) and the known joint distribution of the value of Brownian motion with drift and its running minimum (see Borodin and Salminen (1996)), from

$$\begin{aligned}\mathbb{P}\left(x^-_k < \mathcal{X}_{t_k} < x^+_k, I^{\mathcal{X}}_k > 0 \,|\, \mathcal{X}_{t_{k-1}} = x\right) \\ = \mathbb{P}\left(\hat{Y}_k(x^-_k; x) < y^{(k)}_{t_k} < \hat{Y}_k(x^+_k; x), I^y_k > \hat{Y}_k(0; x)\right)\end{aligned} \tag{3.23}$$

where $I^y_k = \inf_{t_{k-1}<u<t_k} y^{(k)}_u$ and $\hat{Y}_k(\,\cdot\,;\,\cdot\,)$ is defined in (3.20).

Forward Induction and Calibration

For the Markov states underlying a simple lattice implementation, let us consider for each time t_k the possible non-zero levels of the stock prices plus two possible Markov states associated with default, $\mathfrak{D}^1_{t_k}$ and $\mathfrak{D}^2_{t_k}$. The state $\mathfrak{D}^1_{t_k}$ corresponds to default within the interval $(t_{k-1}, t_k]$ and state $\mathfrak{D}^2_{t_k}$ to default within $(0, t_{k-1}]$. We refer by $D_{t_k} := \mathfrak{D}^1_{t_k} \cup \mathfrak{D}^2_{t_k}$ to the aggregate default state, which is absorbing.

The transition probabilities between non-zero stock price levels of consecutive times t_{k-1} and t_k are computed using (3.22) as

$$p\left(S_{t_k}^{(j)}|S_{t_{k-1}}^{(m)}\right) = \mathbb{P}\left(x_{j,k}^{-} < S_{t_k}^{(j)} < x_{j,k}^{+}|S_{t_{k-1}} = S_{t_{k-1}}^{(m)}\right),$$
$$p\left(\mathfrak{D}_{t_k}^{1}|S_{t_{k-1}}^{(m)}\right) = 1 - \sum\nolimits_j p\left(S_{t_k}^{(j)}|S_{t_{k-1}}^{(m)}\right).$$

Trivially, $p(S_0) = 1$ at time $t_0 = 0$. Suppose that transition probabilities from S_0 to each node $S_{t_{k-1}}^{(m)}$ between t_0 and t_{k-1} are already known (denoted by $\mathbf{p}_{k-1}$), and are such that the model values of European options and defaultable zero coupons bond of maturity t_{k-1} are sufficiently close to the market. Then, further probabilities of transition between t_0 and t_k are obtained easily from

$$\mathbf{p}_k = \mathbf{T}_{k-1,k}\mathbf{p}_{k-1} \tag{3.24}$$

with $\mathbf{T}_{k-1,k}$ being the (time-dependent) $(N_k^d + N_k^u + 1) \times (N_{k-1}^d + N_{k-1}^u + 1)$ transition matrix with elements $\left[\mathrm{T}_{k-1,k}\right]_{j,m} = p(S_{t_k}^{(j)}|S_{t_{k-1}}^{(m)})$. This uses transition probabilities for which explicit formulas exists, see (3.19),(3.22), involving calibration parameters q_k, σ_k and a_k. The (unconditional) time zero probability of transition to $\mathfrak{D}_{t_k}^1$ is given by $\mathbf{p}(\mathfrak{D}_{t_k}^1) = \mathbf{d}_k.\mathbf{p}_{k-1}$ where the jth vector coordinate $[\mathbf{d}_k]_j = p(\mathfrak{D}_{t_k}^1|S_{t_{k-1}}^{(j)})$ and the probability of default by t_k is given by

$$p(D_{t_k}) = p(\mathfrak{D}_{t_k}^1) + p(\mathfrak{D}_{t_k}^2) = 1 - \mathbf{p}_k.\mathbf{1}. \tag{3.25}$$

Furthermore, $p(\mathfrak{D}_{t_k}^2|\mathfrak{D}_{t_{k-1}}^j) = 1$ holds for $j = 1, 2$, thus $p(\mathfrak{D}_{t_k}^2) = p(D_{t_k-1})$.

Using risk neutral probabilities p and discounting, the values of continent payment streams stemming from derivative securities, possibly with optimal (Bermudan) exercise optionalities, are computed by usual backward computation. More specifically, for the securities to be used for calibration in the sequel, simple concrete formulas are as follows. Let $\mathbf{S}_n$ denote the vector of time-t_n stock price values strictly greater than zero. The time zero values of European call and put options with expiry t_n and strike K are given by

$$V_C(K, t_n) = \frac{\mathbf{p}_n.(\mathbf{S}_n - K\mathbf{1})^+}{\exp(\int_0^{t_n} r_s ds)} \quad \text{and} \quad V_P(K, t_n) = \frac{\mathbf{p}_n.(K\mathbf{1} - \mathbf{S}_n)^+ + p(D_{t_n})K}{\exp(\int_0^{t_n} r_s ds)}$$

where $\mathbf{1}$ denotes a vector of ones. The value of a defaultable zero coupon bond of maturity t_n is $B_D(0, t_n) = e^{-\int_0^{t_n} r_s ds}(1 - p(D_{t_n}))$. A CDS of maturity t_n, paying spread s over accrual periods Δt for protection against losses $(1 - R)$, with no accrued coupon paid at default, has value $V_{CDS}(0, t_n)$ equal to

$$\sum_{i=1}^{n}\left\{\frac{(1-R)\big(p(D_{t_i}) - p(D_{t_{i-1}})\big)B_D(0, t_{i-1})}{\exp(\int_{t_{i-1}}^{t_i} r_s ds)} - s\Delta t B_D(0, t_i)\right\}, \tag{3.26}$$

where we have used that $\sum_j p(\mathfrak{D}_{t_i}^1|S_{t_i}^{(j)}) + p(D_{t_{i-1}}) = p(D_{t_i})$.

Remark 3.2. The valuation of derivatives has been outlined above. For the later purpose of static hedging according to section 2.2, we attribute cash flows to the derivatives by assuming that any (recovery and option) payments made in default are reinvested until the maturity of the respective instruments. For simplicity, we assume that such reinvestments earn a rate of zero return and that

recovery is constant over time. This fits conveniently with the simple implementation described above. It would be straightforward to deal with non-zero reinvestment returns for payments originating at default, by incorporating further default states in the implementation to keep track of the timing of default.

Calibration Algorithm

Calibration is done by forward induction. During the kth induction step, the parameters σ_k, q_k and a_k are varied to minimise the market/model price error of European out of the money (OTM) calls and puts and the price of defaultable zero coupon bonds of maturity t_k. The calibration at time t_k proceeds along the following three steps: 1) Make suitable first guesses for σ_k, q_k (e.g. $\sigma_k = \sigma_{k-1}$, $q_k = q_{k-1}$ for $k > 1$). 2) Set a_k so that defaultable zero coupon bond prices of maturity t_k are fitted within given error tolerance. 3) Calculate OTM call and put option prices $V_{OTM}(K_j)$ for strikes K_j, and compare to market quotes $V^m_{OTM}(K_j)$. If $\sum_j (V^m_{OTM}(K_j) - V_{OTM}(K_j))^2 > tolerance$ adjust σ_k and q_k and return to step 2) (or terminate if no further improvement is obtained).

Remark 3.3. (i) The calibration minimises the absolute error for OTM options; this ensures greater accuracy in relative error terms for options nearer the ATM strike than those further OTM. Parameters a_k in step 2) were calibrated by Newton's method. For step 3) we used Levenberg-Marquardt's least squares algorithm.

(ii) The procedure above was used successfully for the calibrations in section 5 and delivered positive a_k. Additional constraints on (σ_k, q_k) may be imposed in step 3) to ensure that step 2) can be obtained by $a_k \geq 0$ in general.

4 The Extended Heston Model

This section presents an affine extension of the stochastic volatility model by Heston (1993) to incorporate default risk, which is subsumed by the model of equity and credit risk in Carr and Wu (2006). We refer to it as the Extended Heston Model EH. It permits to price many vanilla derivatives by Fourier transform methods. For the valuation of convertible bonds, we apply PDE techniques.

The Model Framework

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$ be a filtered probability space whose filtration satisfies the usual conditions. The probability $\mathbb{P}$ is the risk neutral measure under which security values, discounted by the riskless savings account $B_t = \exp(\int_0^t r(s)ds)$, are martingales. Unless otherwise specified, expectations are taken under $\mathbb{P}$.

The evolution of the *pre-default* log stock price process $X_t = \log S_t$ (on $\{\tau > t\}$) and its stochastic variance Σ_t are given by

$$dX_t = (r(t) + \lambda(t, \Sigma_t) - \frac{1}{2}\Sigma_t)dt + \sqrt{\Sigma_t}dW^1_t, \tag{4.27}$$

$$d\Sigma_t = \kappa(\theta - \Sigma_t)dt + \sigma_\nu \sqrt{\Sigma_t}dW^2_t, \tag{4.28}$$

with $X_0 = \log S_0, \Sigma_0 > 0$, where W^1 and W^2 are correlated, $\mathcal{F}_t$-adapted, Brownian motions with $d\langle W^1, W^2\rangle_t = \rho dt$. Let N_t be an $\mathcal{F}_t$-adapted doubly stochas-

tic Poisson process[2] that admits the intensity $\lambda_t = \lambda(t, \Sigma(t))$. The time of default of the stock issuer is defined as the (totally inaccessible) stopping time $\tau = \inf\{t : N_t = 1 > 0 \equiv N_0\}$.

The intensity of N_t is taken to be affine

$$\lambda(t, \Sigma(t)) = \alpha(t) + \xi\Sigma(t) \tag{4.29}$$

with $\alpha : [0, \infty) \to \mathbb{R}_+$ being bounded and deterministic, and $\xi \in [0, \infty)$ constant.

Pricing by Transform or PDE methods

At time t, the (pre-default) value of a defaultable derivative security with European style payoff $f(X_T, \Sigma_T)$ at T (if $\{\tau > T\}$) is,

$$\begin{aligned} u(t, x, \nu) &:= \mathbb{E}\Big[e^{-\int_t^T (r(s)+\lambda(s,\Sigma_s))ds} f(X_T, \Sigma_T)|\mathcal{F}_t\Big] \\ &= \mathbb{E}\Big[e^{-\int_t^T (r(s)+\lambda(s,\Sigma_s))ds} f(X_T, \Sigma_T)|X(t) = x, \Sigma(t) = \nu\Big] \end{aligned} \tag{4.30}$$

Under the above assumptions, the model is affine. This permits evaluation of (4.30) for several derivatives by inversion of Fourier transforms, see Carr and Madan (1998); Duffie et al. (2000), which we summarise in the Appendix. Such derivatives include European vanilla options and credit default swaps, but not convertible bonds, whose prices will be computed by solving PDEs numerically described in the following section.

Pricing by PDE methods: Convertible bonds

Let $u(t, x, \nu)$ denote the pre-default value at time t of some defaultable European derivative of fixed maturity T, when $X(t) = x$, $\Sigma(t) = \nu$, in the EH model. Assume that the derivative provides payments at rate $g(t)$ to the investor before default, and pays a constant amount R at default. Supposing sufficient smoothness, u then satisfies the PDE

$$\frac{\partial u}{\partial t} + \mathcal{L}u - (\lambda(t, \nu) + r(t))u = -\lambda(t, \nu)R - g(t) \tag{4.31}$$

on $(0, T) \times \mathcal{D} \subset \mathbb{R}^2$. The infinitesimal generator of the process (X_t, ν_t) is

$$\begin{aligned} \mathcal{L} = &\frac{1}{2}\nu\frac{\partial^2}{\partial x^2} + \rho\sigma_\nu\nu\frac{\partial^2}{\partial x \partial \nu} + \frac{1}{2}\sigma_\nu^2\nu\frac{\partial^2}{\partial \nu^2} \\ &+ \Big[r(t) + \lambda(t, \nu) - \frac{1}{2}\nu\Big]\frac{\partial}{\partial x} + \Big[\kappa(\theta - \nu)\Big]\frac{\partial}{\partial \nu}\,. \end{aligned} \tag{4.32}$$

To compute FD solutions to the PDE (4.31) one needs to provide truncation boundary conditions, in addition to the boundary condition at terminal time given by the payoff functional. We do so below, turning to the pricing of a convertible bond. It has maturity T, and notional of 1 currency unit paying some annualised coupon c that accrues over the period Δ^c years, is convertible on coupon payment dates only (after coupon being payed) to ζ stocks and callable by the issuer on conversion dates only for the price CB_c. The holder of

[2] That is, conditional on knowing $\mathcal{H} := \sigma(\{\lambda_t : t \geq 0\})$, the process N_t is a Poisson process with (known) intensity (λ_t) under the larger filtration $(\mathcal{F}_t \vee \mathcal{H})_{t\geq 0}$.

the bond still has the right to convert after the issuer has exercised the right to call.

Denote the value at t of a T maturity defaultable bond with no conversion features by $DB(t,T)$ and assume that the bond recovers a constant R, equal to a fraction of the bond notional at default. The (pre-default, ex-coupon) convertible bond price $u(t,x,\nu)$ is described by the parabolic PDE

$$0 \quad = \quad \frac{\partial u}{\partial t} + \mathcal{L}u - (\lambda(t,\nu) + r(t))u + \lambda(t,\nu)R \tag{4.33}$$

on $(t_i^c, t_{i+1}^c) \times \mathcal{D}_+$ with $\mathcal{D}_+ = \{(x,\nu) \in \mathbb{R}^2 : \nu > 0\}$, describing the (pre-default) value u between coupon payment dates t_i^c, for $i = 1, ...N^c$. Further, the boundary condition at maturity $t_{N^c}^c = T$ is

$$u(T-,x,\nu) \quad = \quad \max(\zeta e^x, \min(CB_c, DB(T,T)))\,,$$

where $DB(T,T) = 1 + c\Delta^c$, while at intermediate coupon dates t_i^c, $i < N^c$,

$$u(t_i^c-,x,\nu) \quad = \quad c\Delta^c + \max(\zeta e^x, \min(CB_c, u(t_i^c+,x,\nu)))$$

holds. The last equation reflects that the holder, after receiving the coupon, has the option to convert, whereas the issuer might call the convertible bond.

To set suitable truncation boundary conditions for the FD implementation of a numerical PDE solution, we postulate the following boundary behaviour

$$u(t,x,\nu) \to DB(t,T) \;\text{ as } x \to -\infty\,, \qquad \frac{\partial^2 u}{\partial x^2} - \frac{\partial u}{\partial x} \to 0 \text{ as } x \to +\infty\,,$$

recalling that x denotes log-prices. For the ν boundaries, we assume $\frac{\partial^2 u}{\partial \nu^2} \to 0$ as $\nu \uparrow \infty$, and for $\nu \to 0$, a condition for $\frac{\partial u}{\partial \nu}$ follows from taking u to satisfy the PDE (4.33) with $\nu = 0$, i.e. $\frac{\partial u}{\partial t} + \tilde{r}(t)\frac{\partial u}{\partial x} + \kappa\theta\frac{\partial u}{\partial \nu} - \tilde{r}(t)u + \alpha(t)R = 0$ for $\tilde{r}(t) := r(t) + \alpha(t)$ at $\nu = 0$.

The numerical PDE solution is computed by FD methods on some finite truncated domain $\mathcal{D} = \{(x,\nu) : x_L \le x \le x_U, \nu_L \le \nu \le \nu_U\}$, which is to be chosen such that the truncation introduces sufficiently small errors. To compute the numerical solution we have used a predictor-corrector FD method that employs operator splitting from Hout and Welfert (2007).

The Forward Equation

To compute moments or other statistics of the P/L of a statically hedged portfolio in the EH model, see Section 5.4 and Remark 2.3, we need the pricing kernel density, that is the risk neutral (unconditional) transition density with appropriate discounting (killing). It is computed from the forward PDE, that describes the evolution of the density (including the discounting), by FD methods.

Recall the backward equation for the pricing density, that is the suitably discounted (by risk less rate and intensity to default) transition density function $p(t,x,\nu,t',y,\mu)$ from (x,ν) at time t to (y,μ) at t'. Supposing sufficient smoothness of p, the backward Kolmogorov equation is $\frac{\partial}{\partial t}p + \mathcal{A}p = 0$ with operator $\mathcal{A} = \mathcal{L} - (\lambda(t,\nu) + r(t))$ acting on (x,ν) in $\mathcal{D}_+ = \{(x,\nu) \in \mathbb{R}^2 : \nu > 0\}$, for fixed (y,μ). To describe the forward evolution by a forward equation for the pricing

density associated to the process $(X_{t+s}, \Sigma_{t+s})_{s\geq 0}$ starting at $X_t = x, \Sigma_t = \nu$ at t (fixed), one needs the formal adjoint operator

$$\mathcal{A}^{*,t} := \frac{1}{2}\mu\frac{\partial^2}{\partial y^2} + \rho\sigma_\nu\mu\frac{\partial^2}{\partial y\partial\mu} + \frac{1}{2}\sigma_\nu^2\mu\frac{\partial^2}{\partial\mu^2} - \Big[r(t) + \lambda(t,\mu) - \frac{1}{2}\mu - \rho\sigma_\nu\Big]\frac{\partial}{\partial y} - \Big[\kappa(\theta-\mu) - \sigma_\nu^2\Big]\frac{\partial}{\partial\mu} + \kappa - (\lambda(t,\mu) + r(t))$$

that is acting on the forward arguments y and μ of the pricing density $\widetilde{p}(s,y,\mu) := p(t,x,\nu,s,y,\mu)$. The forward equation then reads $-\frac{\partial}{\partial s}\widetilde{p} + \mathcal{A}^{*,T-s}\widetilde{p} = 0$. With suitable initial condition ($\widetilde{p}(0,dy,d\mu) = \delta(x-y,\nu-\mu)$) and (truncation) boundary conditions, the pricing density can be computed numerically using FD methods, mentioned above. In the examples in section 5, the boundary $\mu = 0$ is non-attainable, since $\sigma_\nu^2 \leq 2\kappa\theta$ holds, and $\widetilde{p} = 0$ on $\partial\mathcal{D}_+$ was used as approximate truncation boundary condition.

5 Static Hedging example: Convertible Bond

We demonstrate the static hedging method of this paper by applying it to hedging a convertible bond with European calls and puts, credit default swaps and zero coupon bonds. We hedge a trade called CB1; a 100 notional, 2 year maturity convertible bond with a 5% coupon paid semi-annually. The bond may be converted by the holder for 0.9 shares or called by the issuer for 120 on coupon dates only. If the bond issuer exercises the right to call the bond, the bond holder may still convert.

5.1 DD+ Model Calibration

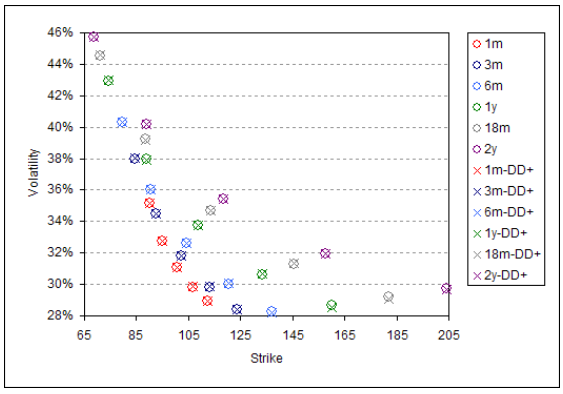

Figure 1: Calibration accuracy (in terms of Black Scholes implied volatility) of the DD+ model to the EH model computed option prices.

We assume that the market values equity and credit derivatives using the EH model with $r = 0.04, \kappa = 1.0, \theta = 0.35^2, \sigma_\nu = 0.3, \rho = -0.6$ and $\xi = 0.35$, $\alpha \equiv 0$. We assume initial values for the stock price and variance of $S_0 = 100$, $\Sigma_0 = 0.09$ (i.e., initial stock volatility of 30%). Recovery for all default sensitive securities is 40% of the notional position prior to default. We assume the trader

Maturity	Intensity	CDS Spread
1m	3.1960%	1.9610%
3m	3.2802%	2.0017%
6m	3.3893%	2.0877%
1y	3.5595%	2.1791%
18m	3.6829%	2.2573%
2y	3.7742%	2.3047%

Table 1: Default intensity and credit default swap spreads calculated using the EH model and assuming 40% recovery in default.

Time t_k	Vol σ_k	Skew q_k	Intensity a_k
0.083	0.3004062577	0.0430075262	0.0319607070
0.25	0.3008774584	0.1397628001	0.0319505343
0.5	0.3050992656	0.2685696661	0.0328532553
1.0	0.3121133276	0.4301136463	0.0337547937
1.5	0.3266560551	0.5250142422	0.0343758047
2.0	0.3404775784	0.5716379599	0.0318235221

Table 2: DD+ model calibration parameters with p = 2 and b_k chosen to match EH model defaultable zero coupon bonds.

hedging CB1 observes the EH market prices of the hedge instruments, but uses and calibrates his DD+ model (with $p = 2$) to reprice hedge instruments as close to the market as possible. For DD+ calibration parameters see table 2.

To simulate the hedger's choice among a (small) given number of liquid options in the market with plausible strikes for respective maturities, we take the given (European) options to have strikes equivalent to those Black Scholes option deltas, where call and put delta are equal or (out of the money) calls respectively puts have deltas of 10% and 25%, for each given expiry dates.

The accuracy of the DD+ model calibration to option prices is illustrated in figure 1 in terms of Black Scholes volatility at each option expiry and strike. Option volatility errors are less than 1% along with perfect repricing of CDS hedge spreads given in table 1.

5.2 The Weak Static Hedge Performance

We consider the hedge to be perfect if the value of CB1 plus all hedging instruments (the hedged portfolio) is zero at all expiry times of the hedge instruments (6m, 1y, 18m and 2y). A perfect hedge is unlikely so we judge the performance of the hedge by considering the distributional properties of the non-discounted profit and loss (P/L) for the hedged portfolio at the CB1 coupon payment dates.

Using the calibrated DD+ model and the optimal static hedge algorithm, a hedge portfolio is computed for hedging the value of the convertible bond CB1 at its coupon dates. The notionals of each hedge instrument are given in table 3. At the time the convertible bond CB1 is issued, the trader holds CB1 plus the portfolio given in table 3. No change to the hedge is made over the lifetime of the bond unless the bond is converted, called or default occurs.

#	Trade	Expiry	Notional	Details
1	Par CDS	2.00	102.36196	Spread 0.023038
2	OPTION	2.00	0.044629	CALL strike 68.82
3	OPTION	2.00	-0.164464	CALL strike 89.03
4	OPTION	2.00	-0.829655	CALL strike 118.53
5	OPTION	2.00	0.064043	CALL strike 157.80
6	OPTION	2.00	-0.017375	CALL strike 204.16
7	ZCB	2.00	-101.4171	
8	Par CDS	1.50	-1.043921	Spread 0.022445
9	OPTION	1.50	0.159955	CALL strike 113.60
10	OPTION	1.50	-0.23741	CALL strike 145.55
11	OPTION	1.50	0.085212	CALL strike 181.92
12	ZCB	1.50	-1.381842	
13	Par CDS	1.00	-1.084435	Spread 0.021725
14	OPTION	1.00	0.025512	CALL strike 74.12
15	OPTION	1.00	-0.10441	CALL strike 88.93
16	OPTION	1.00	0.404295	CALL strike 108.87
17	OPTION	1.00	-0.418216	CALL strike 133.29
18	OPTION	1.00	0.053135	CALL strike 159.91
19	ZCB	1.00	-1.345131	
20	Par CDS	0.50	-1.111014	Spread 0.020525
21	OPTION	0.50	-0.02296	CALL strike 90.43
22	OPTION	0.50	0.108459	CALL strike 104.34
23	OPTION	0.50	0.383026	CALL strike 120.39
24	OPTION	0.50	-0.677842	CALL strike 136.94
25	ZCB	0.50	-1.3163	

Table 3: Hedge portfolio computed using the DD+ for CB1 primary instrument (CDS = credit default swap with recovery 40%, ZCB = zero coupon bond).

Notice that the hedge portfolio only contains a subset of the available hedge instruments at each hedging time. This subset of hedge instruments has been chosen by experimentation (algorithms to select a small number of regressors are discussed in Lamberton and Lapeyre (1993), section 3.1): for the first hedge calculation all hedging instruments were used; this generated P/L statistics[3] in table 4 - Panel A, that we considered to be the 'benchmark' we could achieve. A second hedge calculation with all hedge instruments except puts resulted in minimal worsening of P/L statistics compared to our benchmark. And this second hedge also included calls with a few relatively small hedge notionals. From this hedge portfolio, the smallest notional call was removed and a new hedge portfolio computed; if the P/L statistics were suitably close to our benchmark we left the removed call out of the hedge portfolio. This procedure was repeated until we were left with call hedge positions greater than 0.01 where we stopped. The P/L statistics for this reduced portfolio are given in 4 - Panel B that are clearly close to those of Panel A. Table 4 - Panel C shows more dramatic worsening of P/L statistics for a portfolio that contains credit default swaps, zero coupon bonds but excludes all put options and, most importantly,

[3]P/L statistics are given in currency units.

Time	P/L Mean	Variance	Skew	Kurtosis
Panel A - 'Benchmark' Hedge using all options				
2.0	0.000000	0.441849	-0.7308	4.15437
1.5	0.000000	0.426554	-1.3938	8.93081
1.0	0.000000	0.531097	0.84611	4.42561
0.5	0.000000	0.241966	-0.9791	12.7067
Panel B - Hedge using portfolio in table 3				
2.0	0.000000	0.443657	-0.7267	4.14161
1.5	0.000000	0.432481	-1.2046	8.60745
1.0	0.000000	0.533097	0.84642	4.42283
0.5	0.000000	0.242496	-0.9768	12.6853
Panel C - Hedge without 2 lowest strike calls				
2.00	0.000000	1.214579	1.70296	4.78944
1.50	0.000000	0.523927	-0.3654	5.99832
1.00	0.000000	0.770099	1.22743	4.80507
0.50	0.000000	0.282684	-0.7937	9.97402
Panel D - Only zero coupon bonds				
2.00	0.000000	1346.182	0.92058	6.40640
1.50	0.000000	825.0955	0.34701	6.54585
1.00	0.000000	435.4891	-0.5850	8.27132
0.50	0.000000	165.0104	-2.6315	18.2412

Table 4: A comparison of DD+ model P/L statistics for different hedge portfolios. Panel A P/L statistics are for a CB1 hedge portfolio using CDS, ZCB and all call and put options. Panel B P/L statistics correspond to the portfolio in table 3. Panel C P/L statistics correspond to a hedge portfolio using CDS, ZCB, calls but not the lowest two strikes and no puts.

the two lowest strike call options at all hedge time horizons. Panel D contains P/L statistics when the convertible bond has not been hedged at all, only zero coupon bonds have been used to set the expected P/L to zero. These numbers demonstrate the extent of P/L variance reduction achieved through the hedge, for all cases considered.

Figure 2 shows plots of the DD+ model P/L for the hedged portfolio (CB1 plus hedges) over a range of future stock prices at each CB1 coupon payment date and the DD+ model probability distribution for reaching stock prices at each CB1 coupon date given the initial stock price $S_0 = 100$. We observe from figure 2 that as we step back in time from the 2 year time horizon to today the hedge quality at coupon dates tends to worsen in the region of high stock prices. The probability density graph shows that for each hedge time horizon the regions of high P/L correspond to low likelihood, high stock price values.

The majority of the P/L error is in the jagged region for stock prices of 100 to 130. Our further tests show that if we add a 2 year call option of strike 113.88 (Bond Value/conversion ratio = (100 + 100 x 5% x 0.5)/0.9=113.88) all the 2y hedge options have zero notional except for the 113.88 option and we have a perfect hedge (almost zero P/L for all stock price levels). This is not surprising, given the conversion ratio. However the hedge error at earlier times is mostly

Time	P/L Mean	Variance	Skew	Kurtosis
Panel A - Hedge using portfolio in table 3				
2.0	0.000000	0.443657	-0.7267	4.14161
1.5	0.000000	0.432481	-1.2046	8.60745
1.0	0.000000	0.533097	0.84642	4.42283
0.5	0.000000	0.242496	-0.9768	12.6853
Panel B - Hedge in DD+ model with market price changes				
2.0	-0.00147	0.402973	-0.7818	5.23895
1.5	-0.18047	0.572386	-2.0936	12.7339
1.0	-0.45498	0.696917	-0.6579	3.50667
0.5	-0.68086	0.747691	-2.6296	16.5599
Panel C - Hedge in EH model				
2.0	0.039114	0.507971	-0.523523	3.905897
1.5	-0.213219	0.71434	-1.628541	7.805858
1.0	-0.322333	1.64298	-0.26542	4.644496
0.5	-0.338016	1.338578	-2.071667	10.23964

Table 5: A comparison of hedge P/L statistics due to changes in the DD+ model parameters (Panel B) and a change in the valuation model from the DD+ model to the EH model (Panel C).

unaffected by this perfect 2y hedge and P/L error remains at the levels given in table 4.

5.3 Robustness of the Static Hedge to Parameter Risk

On a daily basis the trader hedging CB1 will recalibrate the DD+ model using observable market option prices and credit default swap spreads and then revalue the hedged portfolio. We demonstrate how the hedge performs as the DD+ model calibration parameters change due to changes in market prices used in the DD+ model calibration.

To generate a *realistic* scenario of changes to the DD+ model calibration parameters we continue with the assumption that the market values equity and credit derivatives using the EH model. We consider the scenario where immediately after CB1 is hedged with the portfolio in table 3 the parameters the market uses in the EH model change resulting in significant changes to all option volatilities and par credit default swap spreads. On observing the market price moves, the DD+ model is re-calibrated to match the new market price data and the hedged portfolio is revalued using the re-calibrated DD+ model. The positions of the hedge instruments are not changed in any way.

We assume that the market EH model has new post-hedge parameters $\theta = 0.36^2$, $\xi = 0.5$ and spot volatility is now 0.35 (originally $0.35^2, 0.35$ and 0.3 respectively) and all other parameters remain unchanged. Using post-hedge EH model parameters in the EH model we see an increase in the market option volatilities of figure 1 ranging from 4.5% to 8% and CDS spreads are nearly double those in table 1. The DD+ model is now re-calibrated to match the changed (post-hedge) option prices and CDS spreads and the hedged portfolio is revalued. The hedged portfolio's P/L statistics associated with this revaluation

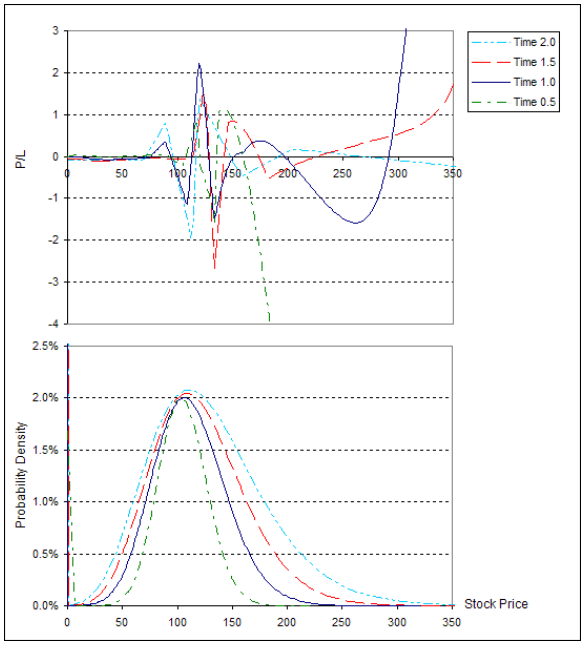

Figure 2: DD+ model hedged portfolio P/L by stock price at coupon dates.

are given in table 5 Panel B. The performance of the hedge after market price moves should be assessed by comparison of table 5 Panel A and Panel B. There is clearly a loss of hedge accuracy as the P/L variance in Panel B has increased over Panel A for the 18m, 12m and 6m hedge time horizons; skew and kurtosis is larger for almost all of Panel B compared with Panel A. However given the relatively dramatic increase in the market prices for our scenario these P/L statistics suggest the hedge may perform well for more typical market moves.

Scenario	Time 0.5	Time 1.0	Time 1.5	Time 2.0
-	0.00000	0.24197	0.53110	0.42655
25% rel skew shift down	0.26903	0.58254	0.43737	0.44592
25% rel skew shift up	0.24933	0.54138	0.40585	0.45411
0.25 abs skew shift down	0.30975	0.62355	0.43774	0.46870
0.25 abs skew shift up	0.28089	0.54040	0.40408	0.45234
25% rel vol shift down	1.51509	2.01173	0.54917	0.57912
25% rel vol shift up	2.00364	1.29474	0.88676	0.38396
25% abs vol shift down	0.83953	2.20812	1.14961	1.38601
25% abs vol shift up	15.54973	9.76510	4.46552	0.24231

Table 6: P/L Variance for scenarios of changes to model parameters in table 2 with other parameters fixed. Top (bottom) rows refer to q_k (σ_k) changes.

Table 6 provides a comparison of P/L variance for different scenarios of changes to the calibrated DD+ model parameters in table 2 with all of other parameters of the model fixed. The first 4 rows are changes to all of the model

skew parameter values q_k, the last 4 rows are to all volatility values σ_k (*rel* means relative, *abs* means absolute). We observe little effect on the hedge performance for all scenarios of skew parameter changes. Volatility scenarios show a larger, but reasonably small effect on the hedge performance for all but the last scenario. Even for this case the statically hedged portfolio demonstrates only a small fraction of the P/L variance of the unhedged portfolio under the original EH model (see table 4 Panel D) and so we might still consider the hedge as successful.

5.4 Robustness of the Static Hedge to Model Risk

Our assumption was that the market values equity and credit derivatives using the EH model. We may demonstrate the robustness of the hedge methodology to model risk by valuing the P/L of the hedged portfolio in the EH model.

To this end, we calculate this P/L over a suitable range of future stock price and stock variance at each CB1 coupon date. To calculate P/L statistics for the EH model we require the transition probability from S_0, Σ_0 at time zero to all values in the same range of future stock price and variance at each CB1 coupon date. The valuation of the hedged portfolio will make use of the inverse Fourier transform methods and the numerical solution of PDEs from Section 4. The required transition probabilities may be calculated by numerical solution of the forward equation given in Section 4.

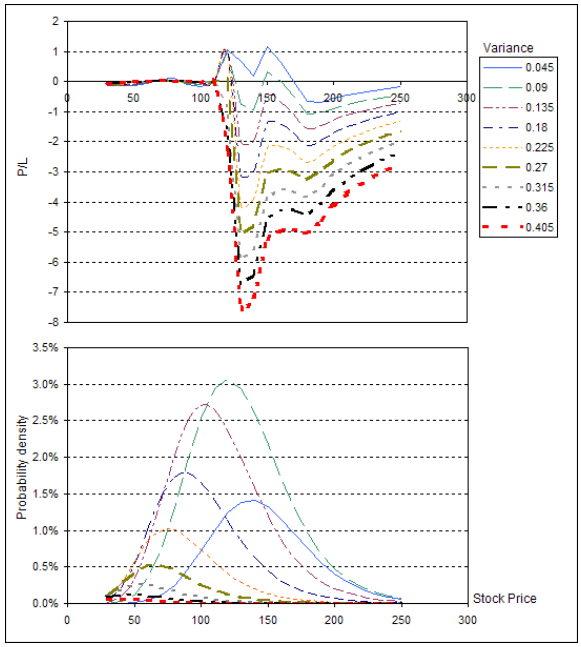

Figure 3: EH model 18 month hedged portfolio P/L and discounted stock price distribution (pricing density) for varying levels of stock variance.

Figure 3 displays the hedged portfolio P/L (CB1 plus hedge portfolio in table 3) against values of the 18 month stock price. Each plot on this graph is for different values of the 18 month stock variance, i.e., values of Σ_t for $t = 1.5$.

Figure 3 also displays the probability of transition from the spot stock price and variance of 100 and 0.09 respectively to 18 month stock price (x axis) and variance (each curve). In general, these plots show that the hedged portfolio P/L is relatively small in a region where the 18 month stock variance is near the initial stock variance (variance 0.09, equivalent to 30% volatility in the DD+ model). Moving into regions of higher variance results in higher P/L however figure 3 shows that these points of higher P/L typically correspond to lower probabilities of transition.

The performance of the hedged portfolio when valued in the EH model is given by table 5 Panel C. A comparison of the P/L variance in table 5 Panel A and Panel C shows that the hedge computed using the DD+ model, then used as a hedge for CB1 in the EH model does not perform as well as for the DD+ model given in table 5 Panel A. However given the significant differences in the two models and that the variance in table 5 Panel C remains low compared to the CB1 notional (below 2% of CB1 notional) we might consider the hedged portfolio in this example a success.

5.5 Weak Static Hedging versus Dynamic Delta Hedging

The performance of static hedging the convertible bond CB1 is now compared with an alternative (ad-hoc) dynamic delta hedging strategy when there are transaction costs (TCs). Let $0 = T_0 < ... < T_M = T$ be a subset of the occurrence times in A1 such that the delta hedging strategy is rebalanced at times T_k, for $k < M$, to hedge a primary instrument with maturity T.

The delta hedging strategy considered for a long CB1 position is as follows. At each rebalancing time T_k, short Δ_k of stock to reduce the sensitivity of the total portfolio to stock price movements between T_k and T_{k+1}. For each hedging period $[T_k, T_{k+1})$ a par CDS with maturity T and semi-annual coupons is held, with notional being set such that the P/L of the hedged portfolio is zero if default were to occur at time T_k. At time T_{k+1} this CDS position is closed (by entering an opposite CDS position). The delta hedge CDS price changes were not hedged as the CDS spreads in the DD+ model were found to be relatively insensitive to stock price changes. The positions in stock, CDS and the primary instrument CB1 are funded from the savings account.

To compare delta hedging with static hedging, paths of the stock price are generated at each rebalancing time with DD+ dynamics given in (3.14) using Monte Carlo simulation. Each simulation path is then used to traverse a DD+ stock price lattice computed from the calibration parameters given in table 2 to compute the value of delta hedged portfolio at each rebalancing time.

The P/L statistics for delta hedging a long CB1 position with monthly and 2 weekly rebalancing are given in the tables below where the results are generated using 5000 simulation paths. The 2nd and 3rd columns of table 7 contain the time-T_k non-discounted value of the P/L mean and variance for delta hedging a long CB1 position up to time T_k with monthly re-balancing (T_k given in the first column). The last two columns contain the mean and variance of the time zero discounted value of TCs for delta hedging up to T_k. We notice that the use of CDS hedges greatly reduces the variance of the hedged portfolio P/L for relatively small TC and small reduction of the hedged portfolio mean P/L. Table 8 contains similar information for delta hedging CB1 with rebalancing every 2 weeks rather than monthly. The increased frequency of the rebalancing

Time	P/L Mean	P/L Variance	TC Mean	TC Variance
Panel A - 1m rebalancing, no CDS used, TC 0bp				
2.0	3.803940	169.7663	0	0
1.5	2.553855	106.9060	0	0
1.0	1.123477	55.16908	0	0
0.5	0.024917	19.15462	0	0
Panel B - 1m rebalancing, no CDS used, TC 50bp				
2.0	2.729340	163.4823	-0.95934	0.333858
1.5	1.701802	102.5294	-0.78250	0.245581
1.0	0.539416	53.41137	-0.55114	0.131598
0.5	-0.30599	19.06527	-0.32077	0.024667
Panel C - 1m rebalancing, CDS used, TC 0bp				
2.0	3.629878	41.22700	0	0
1.5	1.883685	25.01643	0	0
1.0	-0.16951	11.39029	0	0
0.5	-1.85392	2.569395	0	0
Panel D - 1m rebalancing, CDS used, TC 50bp				
2.0	2.282648	38.90189	-1.22478	0.303458
1.5	0.762635	23.21537	-1.03806	0.211020
1.0	-0.98100	10.81715	-0.76211	0.109771
0.5	-2.33317	2.714840	-0.45173	0.020833

Table 7: The P/L and TC from monthly delta hedging, for varying the use of CDS for default risk hedging and the level of TC.

increases the TCs but reduces the variance of P/L for the delta hedged portfolio, as one would expect.

In comparison to the delta hedge, the static hedge of table 3, results in a hedge portfolio P/L variance, given in table 4 panel B, that is significantly lower than the delta hedge with the lowest P/L variance, given in table 8 Panel B. If the TCs for options and CDS in the static hedge of table 7 are computed in the same way as they were for the delta hedge, the expectation of discounted TCs at time zero for the static hedge portfolio is -0.2084. We see that the static hedge up to the 2 year time horizon has TCs that are a fraction of those for delta hedging (this would still hold even if TCs for OTM options in the static hedge are up to almost six times greater) and significantly out performs the delta hedge, in terms of P/L variance minimisation at the coupon dates of CB1. Note that the mean TCs for delta hedging with re-hedging every two weeks are already mostly larger than the standard deviation of the static hedging error. Re-hedging at a higher frequency would decrease the variance of the delta hedging error but would further increase the TCs.

We note that the observed positive mean P/L is mostly due to the long stock option position embedded in the convertible bond, that results in a long gamma. As a caveat, it should also be noted that we have compared an 'ad-hoc' implementation of delta hedging against static hedging here, neither of which having been optimised with respect to TC.

Time	P/L Mean	P/L Variance	TC Mean	TC Variance
Panel A - 2w rebalancing, no CDS used, TC 50bp				
2.0	2.766969	136.6035	-1.23883	0.544749
1.5	1.797473	80.60935	-0.96627	0.375807
1.0	0.453373	45.30512	-0.66135	0.204205
0.5	-0.34726	17.76466	-0.36452	0.039486
Panel B - 2w rebalancing, CDS used, TC 50bp				
2.0	1.862928	30.30983	-1.77443	0.481376
1.5	0.381399	16.64296	-1.47849	0.299304
1.0	-1.29223	8.289043	-1.08026	0.153423
0.5	-2.50896	2.602309	-0.62185	0.029939

Table 8: The P/L and TC from 2 week delta hedging, for varying the use of CDS for default risk hedging.

6 Conclusion & Further Research

We have studied a static minimum variance hedging methodology for hedging equity and credit derivative securities, namely a convertible bond, using other derivatives. The optimal static hedge was shown to be robust under realistic hedge model parameter changes, and also within a qualitatively different 'real world' model. The fact that the hedge performs well under a completely different model suggests that the hedge instruments are well suited for their task, in that they react similarly as the primary instrument to be hedged. This can be related to nice results from Bielecki et al. (2006), showing that a convertible bond (CB) can be decomposed into a pure defaultable bond and an American-type game option. The latter being not traded in practice, this shows a perfect (static) decomposition of the convertible security into ideal but hypothetical derivatives. The game option part permits definition of an implied CB volatility that could be useful for calibration and quoting purposes. In comparison, our paper shows how to compute an approximate static hedge using given vanilla derivatives, after calibrating a model to their prices. The observation that this works rather well is intuitively supported by the insight from Bielecki et al. (2006, 2007) that the CB embeds a floating strike American-type game option.

In general, the method generically depends on the transition probabilities of some underlying Markov process $\mathfrak{X}_{t_k}$, which are usually not determined by, say, European calls and puts alone. If those alone determine the calibration, this can give rise to model risk which should show up in particular with derivatives whose payoffs depend strongly on joint distributional properties of $\mathfrak{X}_{t_k}$. This was demonstrated in Schoutens et al. (2004) or Hirsa et al. (2003) in terms of big variations in the prices of exotics. For the hedge to perform well it is essential to choose appropriate instruments for hedging (calibration), with similar dependencies as the derivative to be hedged. In contrast, statically hedging a forward starting option by European vanilla options maturing at the forward start date would involve high model risk, as an analysis of hedge robustness like in Sections 5.3-5.4 would reveal quickly.

There are interesting avenues for further research and possible extensions:

1. Adaptive calibration. Clearly, a model should be calibrated mainly to

the relevant hedge instruments. Which instruments are relevant? Those that contribute to the hedge! This suggests an adaptive calibration, by *re-calibrating to the hedge*. Supposing an insufficient initial calibration to some market prices (unlike the DD+ example), one could adjust the respective weights in the objective function according to the importance of each option, revealed by the static hedge, to re-calibrate where it matters most. This procedure can be iterated.

2. Hedging of (or with) more exotic options. One could study to what extent other more exotic options could be hedged statically and what the most appropriate hedge instruments are. In reality, options on individual stocks usually are of American type. For simplicity we have used European options but the same algorithm can also deal with American options (of Bermudan type). Since the latter hedge instruments are rather more similar to the options embedded in a convertible bond, the hedge quality should not deteriorate.

3. Alternative models. Let us note that the DD+ model serves just as an example for a model that can be calibrated to given hedge instruments. One can use another model, provided it fits in the framework of Section 2.

4. Parameter risk, as in Section 5.3, may be incorporated into the hedge computation, by computing a (global) minimum variance hedge over an extended probability space with Bayesian prior distribution over parameter scenarios.

Appendix

Pricing by Transforms

For convenience we summarise transform pricing results from Carr and Wu (2006) applied to the EH model. Let

$$\phi(u) = \mathbb{E}\Big[e^{-\xi\int_t^T \Sigma_s ds} e^{iu\widehat{X}_T}\Big|\mathcal{F}_t\Big] = \exp\Big\{iu\int_t^T (r(y)+\alpha(y))dy - A(s) - B(s)\nu\Big\}$$

denote the (analytic) discounted characteristic function of $\widehat{X}_T = \log S_T/S_t = X_T - X_t$ for $T > t$, with $\nu = \Sigma(t)$, $s = T-t$, $A(s) = \frac{\kappa\theta}{\sigma_\nu^2}\Big[2\log\{1 - \frac{(\eta-\kappa^M)}{2\eta}(1 - e^{-\eta s})\} + (\eta-\kappa^M)s\Big]$, $B(s) = [2b_\nu(1-e^{-\eta s})]\,/\,\big[2\eta - (\eta-\kappa^M)(1-e^{-\eta s})\big]$ and with $\kappa^M = \kappa - iu\sigma_\nu\rho$, $\eta = \sqrt{\kappa^M + 2\sigma_\nu^2 b_\nu}$ and $b_\nu = (1-iu)\xi + 1/2(iu+u^2)$.

For a European derivative with payoff $\Pi(k;a,b,\theta,c) = (a + be^{\theta\widehat{X}_T})1_{\{c\widehat{X}_T\le k\}}$ let $P_t(k;a,b,\theta,c) = \mathbb{E}\Big[e^{-\xi\int_t^T \Sigma_s ds}\Pi(k;a,b,\theta,c)\Big|\mathcal{F}_t\Big]$ denote the expectation of the discounted payoff. The (generalised) Fourier transform of P_t equals

$$\mathcal{P}_t(z;a,b,\theta,c) := \int_{-\infty}^{\infty} e^{izk}P_t(k;a,b,\theta,c)dk = \frac{i}{z}(a\phi(zc) + b\phi(zc - i\theta))$$

for z from the domain $\mathcal{C} \subseteq \mathbb{C}$ where $\mathcal{P}_t$ is defined, see Prop. 3 in Carr and Wu (2004). Prices for several derivatives can be calculated numerically by inversion of analytically known Fourier transforms. The value at t of a call option with maturity T and strike K is

$$V_C(K,t,T) = e^{-\int_t^T (r_s+\alpha_s)ds}P_t(-\log K/S_t; -K, S_t, 1, -1)\,.$$

The value at t of a defaultable zero coupon bond with maturity T is $V_{RZ}(t,T) = e^{-\int_t^T (r_s+\alpha_s)ds}\phi(0)$, and $\tilde{q}(t,T) = e^{-\int_t^T \alpha_s ds}\phi(0)$ is the survival probability between t and T. A credit default swap may be valued using the default probabilities and defaultable zero coupon bonds, similarly as in (3.26). The value $V_P(K,t,T)$ at t of a European put option that pays the strike K at expiry T in the event of a default before expiry is given by

$$e^{-\int_t^T r_s ds}\Big[e^{-\int_t^T \alpha_s ds}P_t(\log K/S_t; K, -S_t, 1, 1) + K(1-\tilde{q}(t,T))\Big].$$

References

L. Andersen, J. Andreasen, and D. Eliezer. Static replication of barrier options: Some general results. *Journal of Computational Finance*, 5(2):1–25, 2002.

L. Andersen and D. Buffum. Calibration and implementation of convertible bond models. *Journal of Computational Finance*, 7(2):1460–1559, 2003.

T.R. Bielecki, S. Crépey, M. Jeanblanc, and M. Rutkowski. Arbitrage pricing of defaultabel game options with applications to convertible bonds. To appear in *Quantitative Finance*, 2006.

T.R. Bielecki, S. Crépey, M. Jeanblanc, and M. Rutkowski. Convertible bonds in a defaultable diffusion model. Working Paper, 2007.

A.N. Borodin and P. Salminen. *Handbook of Brownian Motion - Facts and Formulae*. Birkhäuser, Basel, 1996.

N. Bouleau and D. Lamberton. Residual risk and hedging strategies in Markovian markets. *Stochastic Processes and their Applications*, 33:131–150, 1989.

D.T. Breeden and R.H Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651, 1978.

H.M. Brown, D. Hobson, and L.C.G Rogers. Robust hedging of barrier options. *Mathematical Finance*, 11(3):285–314, 2001.

P. Carr, K. Ellis, and V. Gupta. Static hedging of exotic options. *Journal of Finance*, 53(3):1165–1190, 1998.

P. Carr and D. Madan. Option valuation using the Fast Fourier transform. *Journal of Computational Finance*, 2:61–73, 1998.

P. Carr and L. Wu. Time change, Levy processes and option pricing. *Journal of Financial Economics*, 71(1):113–141, 2004.

P. Carr and L. Wu. Stock options and credit default swaps : A joint framework for valuation and estimation. *Working Paper*, 2006.

E. Derman, D. Ergener, and I. Kani. Static options replication. *Journal of Derivatives*, 2(4):78–95, 1995.

D. Duffie, J. Pan, and K. Singleton. Transform analysis and asset pricing for affine jump diffusions. *Econometrica*, 68(6):1343–1376, 2000.

H. Föllmer and M. Schweizer. Hedging by sequential regression: An introduction to the mathematics of options trading. *Astin Bulletin*, 19:29–42, 1990.

H. Föllmer and D. Sondermann. Hedging of non-redundant contingent claims. In W. Hildenbrand and A. Mas-Colell, editors, *Contributions to Mathematical Economics*, pages 205–223. North-Holland, Amsterdam, 1986.

A. Giese and J. Maruhn. Cost-optimal static super-replication of barrier options optimization approach. *Journal of Computational Finance*, 10(3):71–97, 2007.

S. Heston. A closed form solution for options with stochastic volatility with applications to bond & currency options. *The Review of Financial Studies*, 6 (2):327–343, 1993.

A. Hirsa, G. Courtadon, and D. Madan. The effect of model risk on the valuation of barrier options. *Journal of Risk Finance*, 4:47–55, 2003.

K.J. Hout and B.D. Welfert. Stability of ADI schemes applied to convection-diffusion equations with mixed derivative terms. *Applied Numerical Mathematics*, 57:19–35, 2007.

P. Jäckel and R. Rebonato. Valuing American options in the presence of user-defined smiles and time-dependent volatility: Scenario analysis, model stress and lower bound pricing applications. *Journal of Risk*, 4(1), 2001.

M.S. Joshi. *The Concepts and Practice of Mathematical Finance*. Cambridge University Press, Cambridge, 2003.

D. Lamberton and B. Lapeyre. Hedging index options with few assets. *Mathematical Finance*, 3(1):25–41, 1993.

R.C. Merton. *Continuous Time Finance*. Blackwell, Cambridge MA, 1990.

M. Nalholm and R. Poulsen. Static hedging and model risk for barrier options. *Journal of Futures Markets*, 26(5), 2006.

P. Pellizzari. Static hedging of multivariate derivatives. *European Journal of Operational Reserach*, 166:507–519, 2005.

M. Rubinstein. Displaced diffusion option pricing. *The Journal of Finance*, 38 (1):312–217, 1983.

W. Schoutens, E. Simons, and J. Tistaert. A perfect calibration - now what? *Wilmott Magazine*, March, 2004.