TRANCHED CREDIT DERIVATIVES: DIFFERENT OR UNIQUE?
CAN TRANCHED CREDIT DERIVATIVES REALLY BE HEDGED?

R. Rebonato & D.S. Kainth , Royal Bank of Scotland

DESCRIPTION OF THE PRODUCT

Economic rationale for the product. Description of implied correlation (and correlation smile). Analogy with Black implied volatility (smile).

FEATURES OF TCDS

TCDs have features that, taken together, would appear to set them apart from other derivatives products. For instance, their payoffs cannot be perfectly replicated even in principle (and even if trading were continuous and without frictions); volatility does not appear as one of the inputs of one of the most popular models (indeed, of the very model used by traders to communicate prices); they display a stronger-than-usual dependence on correlation; their pay-offs are discontinuous; a popular class of models are non-time homogenous (ie, according to these models the future looks intrinsically different from the present even in the absence of direct information to this effect); the modelling is still in its infancy.

We will argue that these features, while important, have probably received too much attention. In our view, therefore, TCDs are different, but not unique. Modern option pricing rests its conceptual foundations on the theoretical ability to replicate a payoff. When trading in not continuous, and in the presence of transaction costs, however, this replication is far from perfect even if the underlying follows the very stylised dynamics prescribed by the Black and Scholes model. On the other hand, even in situations where perfect replication is not theoretically possible (eg, in the presence of jumps or stochastic volatility) the terminal payoff can be satisfactorily approximated using an imperfect (non-self-financing) hedging strategy[1]. We therefore find that the dichotomous textbook distinction between complete markets, where perfect replication and relative pricing are possible, and incomplete markets, where absolute pricing becomes necessary, is unsatisfactory. We find more useful to think in terms of degrees of replicability, and we therefore concur with Cochrane's (2000) assessment: "…Holding [an] option entails some risk, and the value of that option depends on the 'market price' of that risk - the covariance of the risk with an appropriate discount factor. Nonetheless we would like not to [...] go back to 'absolute' methods that try to price all assets. We can [...] still form an approximate hedge based on [...] a portfolio of basis assets 'closest to' the focus payoff. [..]. Then the uncertainty about the option value is reduced only to figuring out the price of the residual. …".

This is the approach we take in this work. We want to show how well the payoff of a tranched credit derivative can be approximated in practice under a variety of real-world evolutions of credit spreads and default occurrences. We will show that the degree of replicability is satisfactory even when spreads and default occur in ways not allowed by the model. We will therefore treat TCDs on the same footing as other derivatives products, and show that, once looked at this way, they are a certainly a challenging, but in no way a unique product. Before doing so, however, it is useful to discuss at a qualitative level why the 'special' features mentioned above are either not unique, or less intractable than they might at first sight appear.

DEPENDENCE ON CORRELATION

It is certainly true that the value of a TCD depends on correlation. Figure 1, for instance, shows the dependence of the value of a FTD on the correlation for two popular models. Such a dependence,

[1] See, eg, Rebonato (2004) for a detailed discussion of imperfect hedging with stochastic volatility and jump-diffusion processes.

however, is by no means unique and can be found, in order of increasing complexity and impact, in European swaptions, power reverse duals, and all spread options. For the latter product, in particular, recall that, if two rates have the same (time-dependent) volatility, *all* the value of the spread option comes from the correlation. For TCD, correlation (or, more generally, the co-dependence structure) is actually the underlying traded variable, and, as discussed above, can play the same role as the implied volatility in the traditional Black world. Furthermore, all TCD models allow the explicit calculation of the sensitivity of the value to the correlation. This is neither possible nor meaningful whenever one uses the still-popular one-factor interest model.

LACK OF REPLICABILITY

As we mentioned above, we do not believe that perfect or imperfect theoretical replicability is the right metric to look at the risk, and the price, of a derivative. We would like to stress that many situations where it is customary to "pretend" for pricing purposes that replicability is possible are in reality very contrived or unrealistic: in the Black-and-Scholes world, for instance, we must assume that we know with absolute certainty the deterministic volatility (and, if we did, vega hedging would make no theoretical sense). With stochastic volatility it is often claimed that we can complete the market by adding an extra option to our universe of trading instruments. This is only correct, however, if we knew the full process of the option used for hedging, not just its price today. The same applies to jump process when a finite number of jump amplitude ratios exist. As for other modelling approaches, Variance Gamma process, and most non-degenerate Levy process, do not allow perfect replication. Lack of perfect replicability is therefore the norm and not the exception.

ABSENCE OF A VOLATILITY PARAMETER IN ONE OF THE MORE POPULAR MODELS

It is indeed true that the Li model does not include volatility among its parameters. Its use in situations where the volatility of spreads should play a role has therefore been questioned. Recall, however, that the Li model explicitly allows for a finite probability of default at every time after trade inception. This means that the dynamics of the credit spread is assumed to be irrelevant until the moment of default, when the spread can be thought of as instantaneously exploding to infinity[2]. So the Li model would be inappropriate for a product such as a credit spread option (whose value depends on the finite widening of the spread), but can adequately account for products whose payoffs depend on the event of default. Indeed, we show in the following that allowing for drifting, diffusing or even jumping spreads does not hamper the ability to carry out a partial but effective hedge.

DISCONTINUITY OF PAYOFFS

It is notoriously difficult to hedge in practice payoffs that do not depend continuously on an indicator quantity. This state of affairs is commonly encountered in interest rates, equities and, especially, FX options. Tranched credit derivatives fall in this category, but they display one marked advantage, in that the underlying plain-vanilla instruments used for hedging (the credit default swaps) display the same discontinuous payoffs as the TCDs. It can therefore be argued that the discontinuity in the payoff of TCDs is, if anything, easier to hedge than in the case of equity, FX or interest rates.

MATURITY OF MODELLING

It is certainly true that credit products, and TCDs in particular, are relatively new, and that their modelling is rather new. However many of the innovations and insight developed in other contexts are directly transportable and readily applicable. Recall that it was not until the early 1990s that an exogenous set of risk-free bond prices could be recovered by the interest-rate models of the day. Today all the market-implied default times of a basket of risky bonds can be recovered by construction both with structural and with reduced-form approaches. More generally, there are

[2] We are assuming zero recovery. In more realistic cases the argument can be easily modified.

profound similarities between forward rates and hazard rates. It is therefore fair to say that credit modelling might still be in an evolutionary state, but has enjoyed a twenty-year head start.

HEDGING SIMULATIONS: ADDRESSING THE CONCERNS

The key risk factors for TCDs are spread movements and default events. In light of the discussion above, one possible procedure to understand the profit/loss implications of an imperfect trading strategy (such as that obtained by pricing and hedging within the Li copula model) is to simulate the actual trading process and to examine the P/L contingent on simulated credit spread and default behaviour. We are then interested in seeing whether our hedging strategy reduces the variance of our profit/loss profile and the magnitude of the mean return. We show below an initial analysis of the hedging of TCDs by examining in detail the most elementary case, that of first to default baskets. For this case the key risk factor is probably the movements of the spreads; in what follows we will examine the effects on P/L as a result of the stochasticity of spreads – a factor not explicitly included in the Li model.

Even this task presents a significant computational challenge: calculating the price and accurate deltas of these products using a naïve Monte Carlo[3] even for a single time horizon can be very time consuming, which would in turn render running a hedging simulation in a realistic time frame almost computationally unfeasible. We have, however, used methods developed within the group (see Joshi and Kainth, Quantitative Finance) to speed up the computation of prices and deltas by several orders of magnitude – thereby enabling us to carry out the following hedging simulations of these products.

THE HEDGING SIMULATION ALGORITHM.

Our hedging simulation is prefaced on algorithms for producing the real world evolutions of credit spreads. Developing and estimating an appropriate process for the modelling of real world and risk neutral spreads is still an active area of research; we have developed two significantly different, albeit relatively elementary alternatives:

- The ``simple'' model: here, we assume that the instantaneous credit spreads of a firm follow a mean reverting log normal random walk with jumps:

  $$dS_t = \alpha(\beta - S_t)dt + \sigma S_t dW_t + JS_t dP_t$$

  The remainder of the credit spread term structure is assumed to move in parallel with the instantaneous spread. We have allowed for the possibility that the spreads of firms can jump. Within this model, the jump distribution of the spreads is assumed to be uniform, with the caveat that the spreads must remain positive. Spread moves between different assets are taken to be correlated; this correlation is assumed to be identical to the default time correlation.

- We have also put together a more complicated ratings based model. This model is far more realistic, incorporating effects such as changes in ratings, the instantaneous spreads are assumed to follow a CEV process and a Jarrow, Lando and Turnbull model for the spread term structure. The results from such a model are richer; however, the results obtained are qualitatively similar to those obtained on the basis of our simple model. Since our aim is principally expository we will concentrate exclusively on the simpler model described above.

THE HEDGING ALGORITHM.

For the sake of clarity, we describe the steps involved in the hedging algorithm. We assume that we are long a first to default basket bought at par i.e., the value of the spreads received match the default

[3] It is possible to price and compute sensitivities of these products using semi analytical approaches as described by Laurent and Gregory assuming a factor correlation. This method, although faster than Monte Carlo lacks the necessary generality to deal easily with market specified recovery rates and correlations.

protection. In the event of no defaults, this spread is always paid as one large payment at the end of the life of the product. If there is a default, we assume that the spread paid is in proportion to the survival time of the product compared to the original deal maturity.
For the hedges, we assume that the vanilla default swaps pay out annually; futhermore we assume that the payment for these is made ``upfront" - rather like an insurance premium.

Assume that we are hedging a first to default swap on two assets with value P_{CDS}. We can express this value in terms of the prices of the vanilla default swaps, $V_1(\lambda_1)$and $V_2(\lambda_2)$, as $P_{CDS}(V_1(\lambda_1),V_2(\lambda_2))$. We assume that if spreads are at (λ_1,λ_2) then we can hedge ourselves by holding

$$\left(\frac{\partial P_{CDS}}{\partial V_1},\frac{\partial P_{CDS}}{\partial V_2}\right) \approx \left(-\frac{\partial P_{CDS}}{\partial \lambda_1}/-\frac{\partial V_1}{\partial \lambda_1},-\frac{\partial P_{CDS}}{\partial \lambda_2}/-\frac{\partial V_1}{\partial \lambda_2}\right)$$

of the single name default swaps. This hedge ignores Ito terms. We ignore the effects of transaction costs within our model for the sake of simplicity. Several features are apparent: we are parameter rather than delta hedging in the true Black Scholes sense. At the grossest level, this is because we cannot replicate the payoff of a basket default swap using vanilla default swaps. Our overall aim is to demonstrate the effectiveness of the parameter hedging strategy in reducing the variability of the profit/loss.

To gain some intuition as to the general behaviour of the hedgee, the hedges and the residual cash we have plotted these as a function of time through the hedging simulation in Figure 1. Again we assume that we are long a first to default basket i.e., we have bought protection and so we will pay out spreads. As we progress through time, we expect that the value of the product will decline because we have the same liability at the end (i.e., the spread payments) but we have protection for a shorter period of time. We have hedged ourselves by shorting credit default swaps. These are bought at par: hence buying/selling them has no effect on our immediate mark to market. However, we will receive spreads on these, as evidenced by the jump in the mark to market at the one year point. Over the maturity of the deal, our outgoings very nearly match the amount taken in, although there is a slight discrepancy.

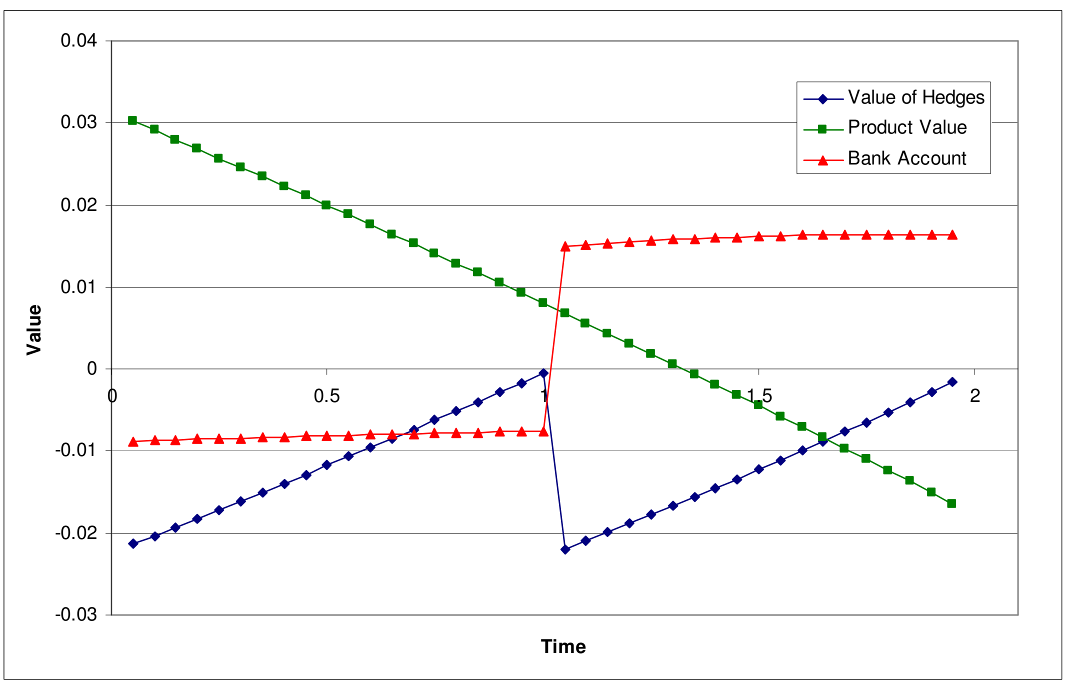

Figure 1: Plots of the values of the hedges, the hedgee and the residual cash through the course of a hedging simulation, where the spreads remain constant through time. The case shown is a first to default on a homogeneous basket of 5 names. Note the discontinuity in the 1 year point in the bank account balance as we receive an upfront spread.

We have analysed the efficacy of our hedging strategy in reducing the P/L given spread movements produced by our `simple' hedging simulator. To better understand the results we have examined the diffusive, drift and jumpy components of the spread behaviours separately. Typical paths which illustrate examples of these behaviours are shown in Figure 2.

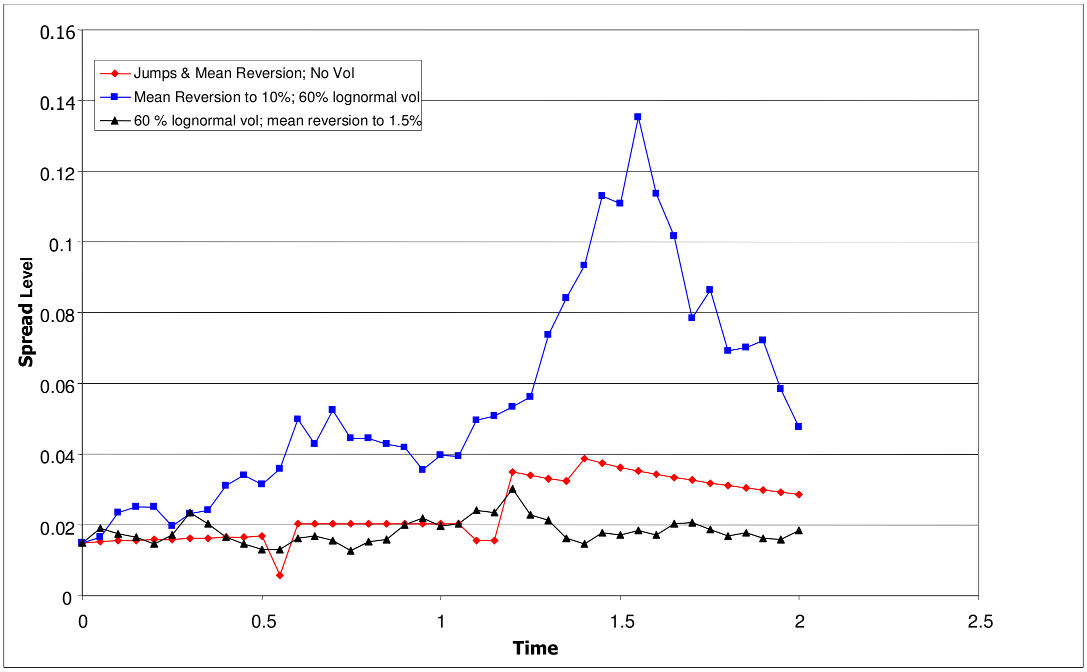

Figure 2: Sample evolutions of spreads for diffusive, mean reverting and jumpy random walks. These sorts of evolutions are ``typical'' of the spread movements produced by our hedging simulator.

For the cases shown below, we have analysed a two year, five asset first to default basket. We assume that the initial spread of all of the obligors is 2%; the recovery rates for all the obligors is set at 0.43. Base correlation between spread movements, jumps and the implied defaults correlation is taken to be 0.3 unless otherwise specified. Interest rates are assumed nonstochastic and are set at 0. In this case the price of the protection is 744bps. For the majority of the results presented below, we have run several hundred (typically between 400 and 1000) paths; the profits/losses generated are typically displayed as a histogram. We have used kernel smoothing techniques to give us smoother, unbiased density plots. Generally we will analyse paths for which no defaults happen and those for which one default happens separately.

Let us begin by thinking about the unhedged case. In this case we expect to pay out 744bps if there is no default (which is roughly 80% of the time) and receive 5700 bps (i.e., 1-R) given that there is a default (the remaining 20% of the time). There is unsurprisingly a huge variation in our P/L profile. Assuming now that we are hedging, we will obtain a P/L graph analogous to the one shown in Figure 3. In this case we assume that the instantaneous spreads follow a lognormal random walk with 60% volatility. It is clear that in the majority of cases (the cases where no default happens) we still lose money – but a much smaller level ~ 60 bps. The variation of the P/L about this mean is also small. When a default happens we again make some money albeit a smaller amount to the unhedged case and with a significant variance. Nevertheless the key message to note is that by hedging we have substantially reduced the variance of our final P/L. We have not included the margin charged up front for this product in our calculations.

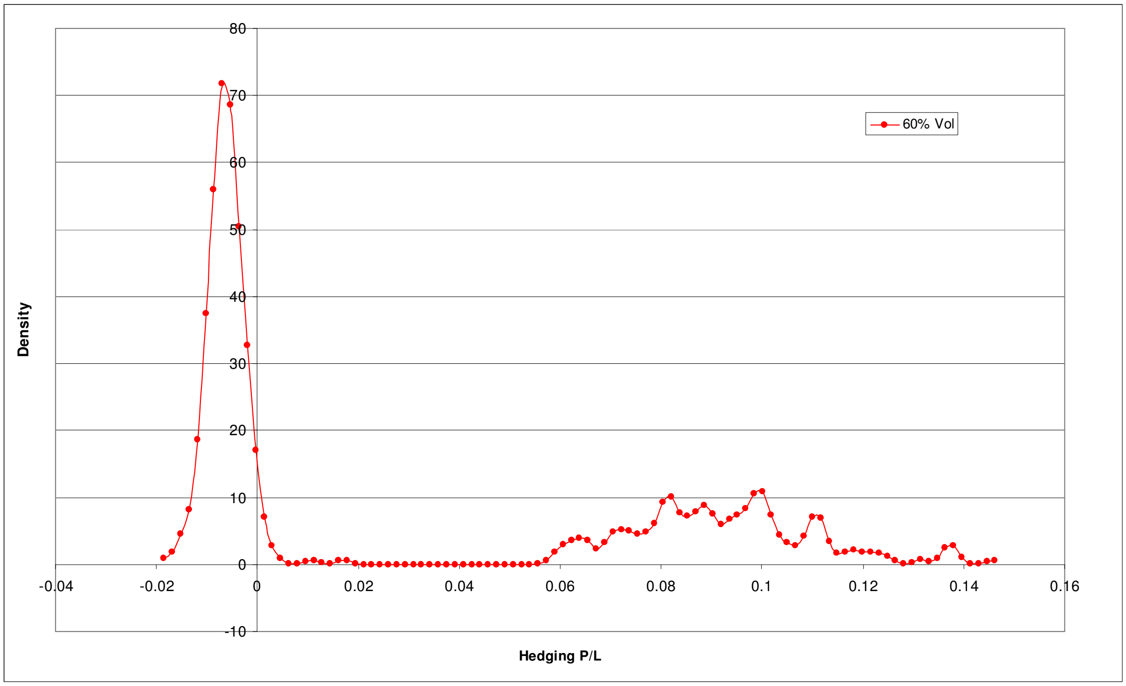

Figure 3: A typical P/L profile from our hedging simulation. The case illustrated assumes a first to default basket of 5 assets, the spreads of which are assumed to follow a lognormal process with no drift and 60% volatility.

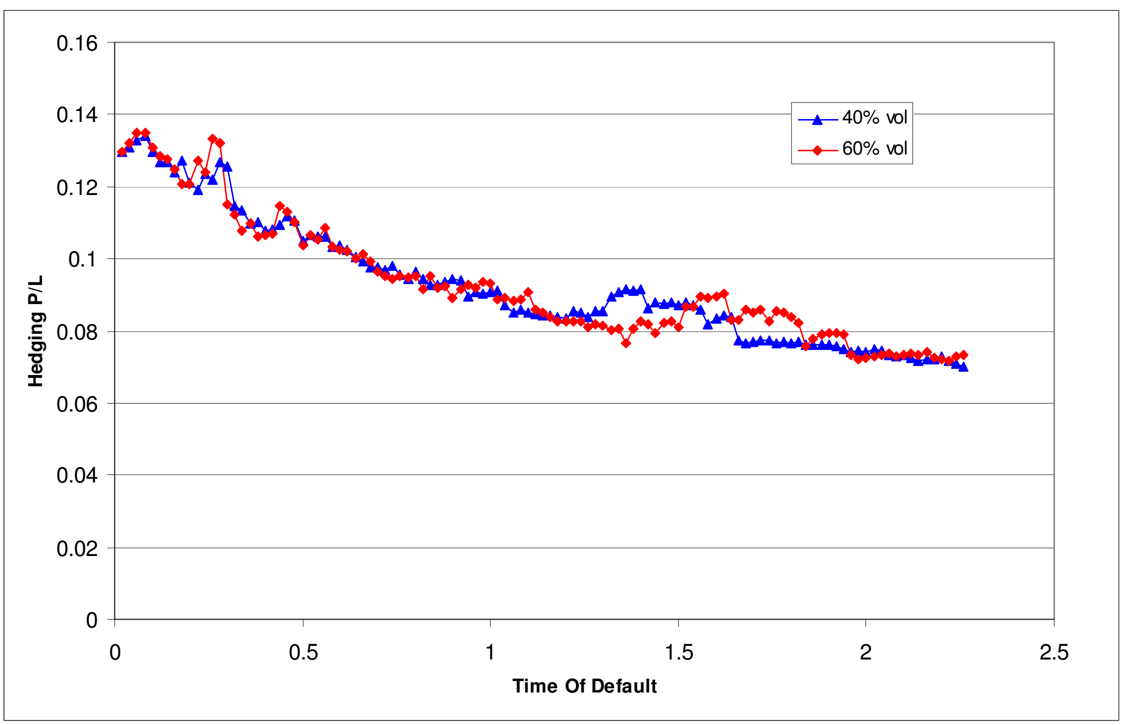

Figure 4: The profit on a first to default basket given a default depends on the time of default and the jump in the spreads of the remaining assets. Here we have for the sake of clarity assumed that there is no jump in spreads following default - we are simply looking at the variability in the profit as a function of default time.

Let us analyse further the P/L variation given that there is a default. In Figure 4 we show the profit on a first to default basket given a default. It is clear that the net gain in cash depends principally on the time of default: if default occurs almost immediately we make ~ 130 bps; this figure declines such that if default occurs at ~ 2 years we make ~ 70 bps only. Increasing spread volatility makes very little difference to these numbers. Here we have assumed that there is no jump in spreads following default (i.e., we are discounting the possibility of default contagion) - we are simply looking at the variability in the profit as a function of default time. Other analysis has demonstrated that we need a pretty substantive jump in spreads before we see that the basket losing money on a default event.

PROFIT/LOSS AS A RESULT OF DRIFT IN SPREADS.

We begin by examining the P/L of the hedging strategy assuming that the spreads are increasing - reverting to a level some way above the one that they are starting at, with a small amount of volatility. The relevant results are shown in **Error! Reference source not found.** where we compare the variability in income as a result of spreads drifting from ~2% up to 12% over the space of two years for two different hedging frequencies; Figure 6 is a tabulation of the means and standard deviations of the P/L for different spread reversion levels.

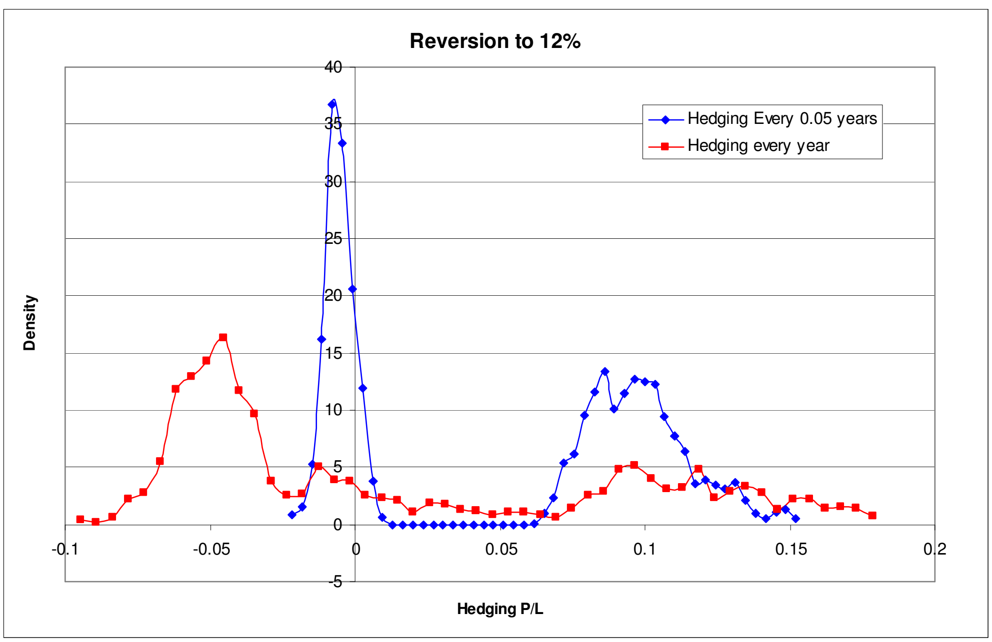

Figure 5: A comparison between the variability in income as a result of spreads drifting up to a mean of 12% over a period of two years. There are two different hedging frequencies illustrated here: one with a rehedge every 3 weeks or so and the other extreme, one hedge every year.

Empirical spreads of firms do exhibit this form of behaviour, where they will drift higher over a sustained period of time. The levels to which we allow the spreads to drift to are not extreme: a number of significant firms post WorldCom and Sept. 11 have had spreads as high as 0.12. The key thing to notice about the two graphs is that if we do not hedge partcularly frequently - we are admittedly comparing something of an extreme, hedging only once a year, with a far more likely hedging frequency - then in the cases where no defaults occur we lose a significant amount of money and perhaps more importantly the variability of our income is high. This is a strong advertisement for the hedging strategy chosen: we are definitely able to reduce the variability in the P/L in the no defaults case.

Reversion Level	Mean (No defaults)	S.D. (No defaults)	Mean (1 default)	S.D. (1 default)
0.02	-0.004	0.0006	0.097	0.0196
0.04	-0.006	0.0014	0.099	0.0175
0.08	-0.007	0.0031	0.1	0.0175
0.12	-0.007	0.0039	0.092	0.0343

Figure 6: Assuming that we hedge every 0.05 years, the mean and standard deviations for P/L assuming no defaults and 1 default.

If we assume that we are hedging every twentieth of a year, Figure 6 reinforces the point that the variability in income, as measured by the standard deviation of the P/L, conditional on no defaults is low ~ 40 bps in our case. The variability in income given that a default happens remains relatively high. This is because the key driver, as discussed above, is in the timing of the default.

Profit/Losses as a result of Volatility.

The copula model of Li does not explicitly incorporate the volatility of the spreads. It is obviously of some concern to see whether or not the pricing and the hedging strategy, as dictated by the copula model can adequately hedge volatile spread movements. To do this we have carried out an analysis similar to that described above only here we assume that the spreads are weakly mean reverting, with no jumps; we increase the lognormal volatility of the spreads to see the effects on the variance of the hedged portfolio.

The results for a variety of cases are detailed in figures 8 and 9 below; sample histograms are shown in fig. 7 for the distribution conditional on no default occurring and in fig. 3 for the whole trade.

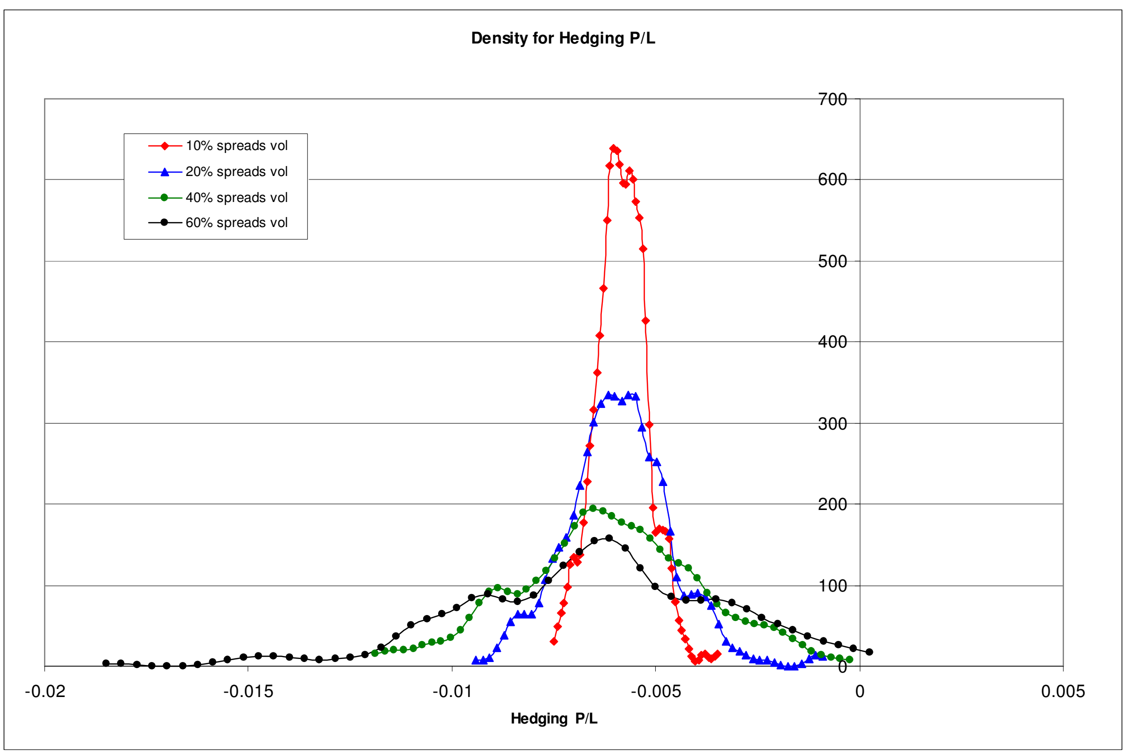

Figure 7: A plot of the density of the profit/loss generated by our hedging strategy, for a first to default basket. We illustrate here cases where there are no defaults. The spreads in this case were taken to be weakly mean reverting, with no jumps and we show results for lognormal vols of 10\%, 20%,40% and 60%.

Volatility	No Defaults (Mean)	No Defaults (S.D.)	One Default (Mean)	One Default (S.D.)
0.1	-0.0058	0.0007	0.0939	0.0137
0.2	-0.0059	0.0013	0.0940	0.0143
0.4	-0.0061	0.0023	0.0919	0.0196
0.6	-0.0065	0.0032	0.0892	0.0245

Figure 8: Assuming that we hedge every 0.05 years, the mean and standard deviations for P/L assuming no defaults and one default as a function of increasing volatility.

Volatility	No Defaults (Mean)	No Defaults (S.D.)	One Default (Mean)	One Default (S.D.)
0.1	-0.0060	0.0006	0.0939	0.0125
0.2	-0.0060	0.0012	0.0945	0.0147
0.4	-0.0076	0.0022	0.0913	0.0220
0.6	-0.0068	0.0032	0.0889	0.0248

Figure 9: As per the table above, except that we now assuming that we hedge every 0.1 years.

We will once again primarily be interested in hedging performance given that there is no default. The two tables above show the mean and variance conditional on number of defaults assuming that we re-hedge every 0.05 years and every 0.1 years respectively. The mean P/L remains pretty much unchanged, regardless of volatility and rehedging frequency. However, it is quite apparent that conditional on no defaults, although the variance of the P/L increases as we increase volatility the numbers are very small: we see a standard deviation of ~30 basis points with volatility at ~60%. The volatilities quoted are lognormal volatilities; this figure represents a typical market volatility.

The standard deviation conditional on one default is as before large; assuming that there is no jump in spreads, the amount we make is usually positive. Once again the large variance in P/L is as a result of the dependence of income on time of default. The key message, nevertheless, is that by parameter hedging within the Li copula model – a model which does not explicitly allow for stochastic spreads – we are able to moderate the variance of our P/L significantly.

BEHAVIOUR OF THE NTH DEFAULT SWAP GIVEN A JUMP UP IN SPREADS.

We have analysed the behaviour of the product and its hedges given the occurrence of a jump in spreads for a first to default swap.

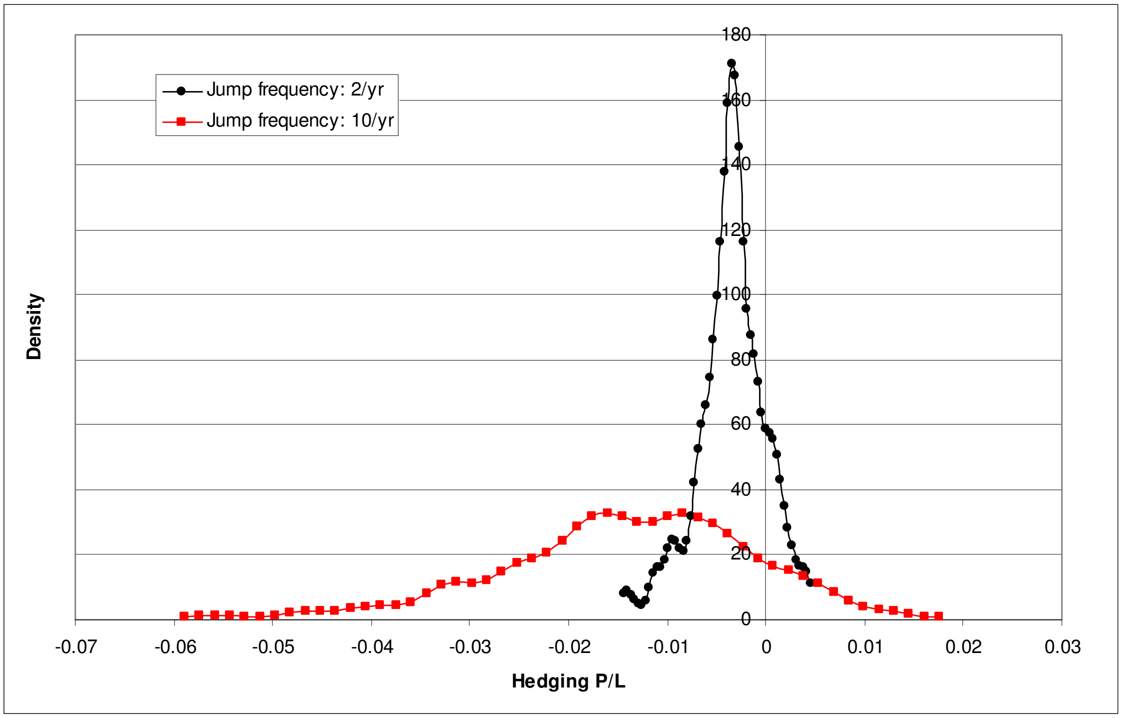

Figure 10: Plot of the density of the profit/loss generated by our hedging strategy, for a first to default basket. We illustrate here cases where there are no defaults. The spreads in this case were taken to be weakly mean reverting, with jumps drawn from a uniform distribution. We have increased the jump frequency; the two cases shown here are for 2 jumps per year and 10 jumps per year. The jump distribution is assumed to be uniform such that the spread can jump to [0.2×CurrentSpread, 3×CurrentSpread] reflecting our belief that most jumps will be upwards and that spreads must remain strictly positive.

As above, we have analysed the effects of jumps on the P/L of the hedging strategy. Corporate spreads do tend to be prone to jumps: there are several explanations for this. Most hinge on the fact that there are a number of corporate defaults over the course of a year; in general these will affect the spreads of other firms --- through real world interactions, such as for example, a failure to pay back loans etc on behalf of one company will lead to a worsening in situations for other firms that trade with this one and so on. Consequently we have varied the jump frequency in our hedging analysis, assuming that the magnitudes and the distributions of the jumps remain unchanged. The jump distribution, as stated above has been assumed to be uniform, conditioned on the current level and such that spreads always remain positive. This is, it would seem, a reasonable first guess at modelling the issues.

Jump Frequency	No Defaults (mean P/L)	No defaults (S.D.)
2	-0.0035	0.0035
4	-0.0046	0.0061
6	-0.0071	0.0082
10	-0.0141	0.0129

Figure 11: Standard deviations of the P/L of the hedging assuming that the spreads undergo jumps. We have increased the jump frequency; the cases shown here are for 2 jumps per year through to 10 jumps per year. The jump distribution is assumed to be uniform such that the spread can jump to [0.2 ×

CurrentSpread, 3×CurrentSpread]. Note that the amount lost and the variation in the final profit and loss conditional on no defaults increases markedly (far more so than for e.g., increasing log normal volatility) as the jump frequency increases.

The results of this set of simulations are presented in Figure 11. It is apparent that the hedging strategy (rather unsurprisingly) is less able to cope with jumps. However, the jump frequencies that we have fed in ~10 per year are somewhat extreme.