Advanced Counterparty Risk and CVA via Stochastic Volatility

Ignacio Ruiz*, Ricardo Pachón†

October 2012

Version 1.2.1

Forthcoming publication in the Journal of Financial Transformation

Abstract

Exposure models in the context of counterparty risk have become central to financial institutions. They are a main driver of CVA pricing, capital calculation and risk management. It is general practice in the industry to use constant-volatility normal or log-normal models for it. Ignacio Ruiz and Ricardo Pachón explain some of the strong limitations of those models and show how stochastic volatility can improve the situation substantially. This is shown with illustrative examples that tackle day-to-day problems that practitioners face. Using a coupled Black-Karasinski model for the volatilty and a GBM model for the spot as an example, it is shown how stochastic volatility models can provide tangible benefits by improving netting effects, CVA pricing accuracy, regulatory capital calculation, initial margin calculations and quality of exposure management.

Counterparty risk has now become a central focus of new developments in the financial industry. This mainly affects three areas in financial institutions: CVA pricing, capital calculation and risk management. The most widely used risk metrics are PFE_t, as a percentile-like measure of future portfolio value, and EPE_t, as the average of the future possible portfolio values after zero-flooring.

*Founding Director, iRuiz Consulting, London. Ignacio is a contractor and independent consultant in quantitative risk analytics and CVA. Prior to this, he was the head strategist for Counterparty Risk, exposure measurement, at Credit Suisse, and Head of Market and Counterparty Risk Methodology for equity at BNP Paribas. Contact: ignacio@iruizconsulting.com

†Risk Model Validation, Credit Suisse, London. This work has been done in the context of a personal collaboration with Ignacio Ruiz. Ricardo is in no way linked to iRuiz Consulting. The content of this paper reflects the views of the author, it does not represent his employer's views. Contact: ricardo.pachon@credit-suisse.com

A key aspect of counterparty risk measurement is the modelling of the underlying risk factors. The most widely used models tend to have normal or log-normal behaviour[1]. These models are popular, arguably, because they usually have analytical solutions and the cost-benefit of implementation is very positive. They are a very good starting point but they also show a number of limitations, which include:

1. Lack of netting with volatility

   In a simple normal model, the volatility is a constant parameter. For that reason, portfolios with volatility-sensitive products, like variance swaps or options, cannot enjoy the full proper netting benefits under that modelling framework.

2. Suboptimal model reactivity

   Counterparty risk models tend to be calibrated to the real measure for capital and risk management purposes, and to the risk-neutral measure for CVA pricing. In either case, given that normal models have a constant volatility, a key challenge that practitioners face is what volatility to calibrate to: 1-month, 3-month, 1-year, 5-year, 10-year?

   On the one hand, if a model is calibrated to short-term volatility, long-term risks can be miss-interpreted and, also, the risk metric (e.g., EPE or PFE) of long-term contracts can be too volatile from day to day. On the other hand, if calibrated to long-term volatility, risk metrics for short-term contracts may not reflect the true risks. This problem is specially acute in crisis periods.

   Finding a right balance between these two poles is very important for a number of key functions within a bank, like the calculation of initial margin, exposure management within the set limits, capital calculations and CVA pricing. As a result of the overly simplicity of normal models, they tend to be calibrated to "somewhere in the middle": typically to a volatility between three to seven years, depending on the portfolio composition and bank policies. Hence risk metrics tend to be hardly reactive to swinging short-term market conditions, as they ideally should be.

   Needless to say, this is suboptimal..

3. Lack of fat tails and skew

   Normal models deliver no skew and no excess kurtosis. However, historical time series indeed show those characteristics. This is particularly relevant for the risk management of exotic trades, that can be very sensitive to the tails of the underlying risk factor distributions.

In this paper we are going to see how a stochastic volatility framework can solve these problems.

[1] From this point, we will refer to these models as "normal" models, for the sake of simplicity.

A Joint Spot-Volatility Stochastic Model for Counterparty Risk and CVA

We find that, often, quantitative publications lack intuitive explanations that practitioner can easiliy relate to. So, first we are going to revise how stochasticity in the volatility affects the diffusion of the spot, in an intuitive way.

Gaining model intuition

The introduction of a stochastic volatility in a normal or log-normal process alters the distribution of the spot during the simulation. To gain some insight, let's say for a moment that the spot (X) we are diffusing follows a normal process

$$dX_t = \sigma_t \, dW_t^X \tag{1}$$

and that the volatility follows a stochastic normal process too[2], so that

$$d\sigma_t = \eta \, dW_t^\sigma, \tag{2}$$

where dW_t^X and dW_t^σ are Brownian motions. Let's say that both Brownian motions have a correlation structure so that

$$dW_t^\sigma = \rho \, dW_t^X + \sqrt{1-\rho^2} \, dW_t^*, \tag{3}$$

where dW^* is also a Brownian motion and where dW_t^X and dW^* are independent.

Putting equations 1, 2 and 3 together, we find that

$$X_t = \widetilde{X_t} + \eta \, \rho \, \frac{1}{2} \left({W_t^X}^2 - t \right) + \eta \, \sqrt{1-\rho^2} \int_0^t W_u^* \, dW_u^x, \tag{4}$$

where $\widetilde{X_t}$ would be the value of X_t if the volatility were a constant equal to σ_0.

Equation 4 illustrates how adding stochasticity to the volatility we are producing non-normal distributions for the spot. Skew will arise from the chi-squared term ${W_t^X}^2$, and kurtosis will arise from that chi-squared and the $\int_0^t W_u^* \, dW_u^x$ terms.

Let's look at this from a numerical point of view now. As expected from equation 4, the correlation ρ delivers an interesting effect. Figure 1, top panel, shows the skew of the one-year moves for a spot simulation for different values of the correlation. We observe how the skew changes as the correlation moves from negative to positive values.

[2] This model choice is done for simplification in the context of this section.

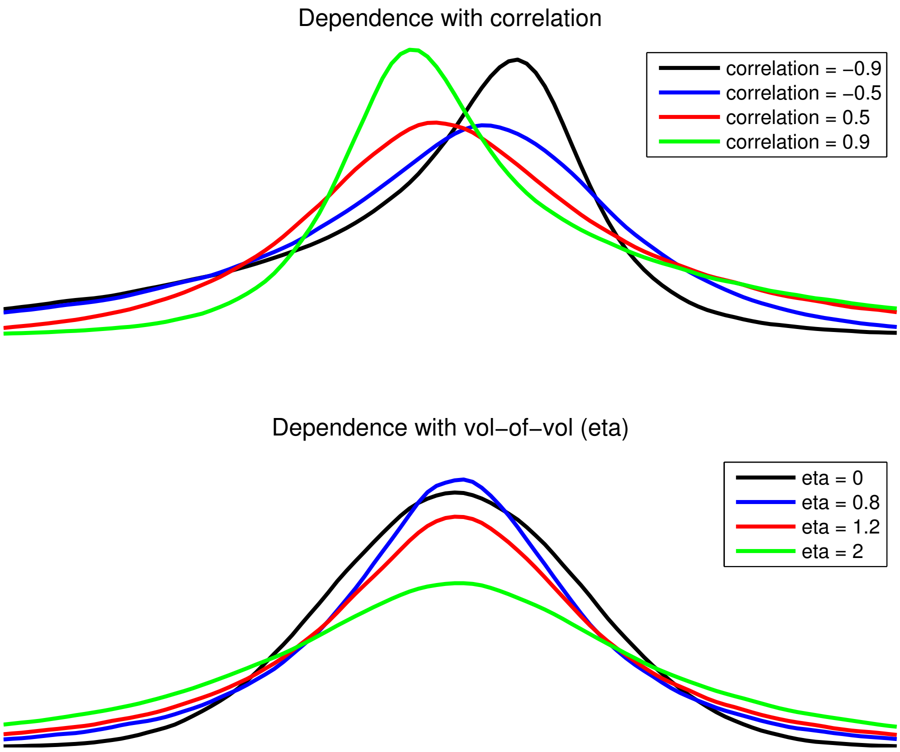

Figure 1: Illustration of dependence of the spot distribution with the correlation ρ and the vol-of-vol η.

The value of the vol-of-vol η also has a notable effect on the spot diffusion. This is depicted in Figure 1, bottom panel. The distribution of the spot process gets leptokurtic as η increases and tends to the normal distribution when it decreases. This is because high values of η result in higher stochastic variation of the volatility process. As a result, high kurtosis is induced because a high η introduces uncertainty in the volatility, which subsequently delivers a "student-T" effect in the spot. We say this because a student-T distribution is precisely created when there is uncertainty around the volatility [9].

We have seen, with a number of simple explanations, how "plugging" stochastic volatility in a normal process we are introducing non-normal effects. It is very important for risk practitioners that may consider implementing this kind of models to understand the impact they are introducing in the spot diffusion.

A Bank Can Benefit Substantially from Stochastic Volatility

There are several stochastic volatility models in the academic literature[7, 4, 6, 8, 5, 2]. However, they tend to gravitate around the subject of deriviative pricing and, to the authors' knowledge, these models have not been studied in the context of the practice of counterparty risk and CVA pricing in financial institutions.

In order to illustrate the power they bring to this space, we wanted to choose a model to show practical results that the practitioner can relate to. In the previous section we chose a normal model both for the spot and the diffusion, as the purpose was to provide intuition as to how stochastic volatility changes the spot diffusion, and that simple model was ideal for that. However, when we come to the real world, that model is not good enough in general[3]. For practical and more realistic examples, we are using the following model: if X is the spot value to simulate, then

$$\begin{aligned} &dX_t = X_t\,\sigma_t\,dW_t^X \\ &d(\ln\sigma_t) = \theta\left(\ln\mu - \ln\sigma_t\right)\,dt + \eta\,dW_t^\sigma \\ &dW_t^X\,dW_t^\sigma = \rho\,dt \end{aligned} \tag{5}$$

where dW_t^X and dW_t^σ are Brownian motions, σ_t is the process volatility and θ, μ, η and ρ are model parameters that will be discussed in the next sections. That is, we choose a Black-Karasinski (BK) model for the volatility and a quasi Geometric Brownian Motion model for the spot[4].

The purpose of this paper is not to provide a survey of all the different stochastic volatility models, and which is good for this or that asset class. However, we think this model is ideal for the needed illustrative purposes because, firstly, volatility is a strongly mean reverting risk factor (see Figure 2) and, secondly, there is some evidence in the literature that it can back-test well [10][5].

Netting Benefit

Typically, in current financial institutions, if risk is measured with a normal or log-normal process, the volatility risk is either missed or "added" on top without any netting benefit. By contrast, with a spot-volatility diffusion model, the correlation between the volatility and the spot is accounted from *within* the same model, and so the institution can benefit from netting effects in an optimal way.

[3] For example, it gives negative volatilities, and is not mean reverting in the volatility.

[4] "quasi" beacause the volatility is not constant, but stochastic.

[5] That study was done in the context of equities.

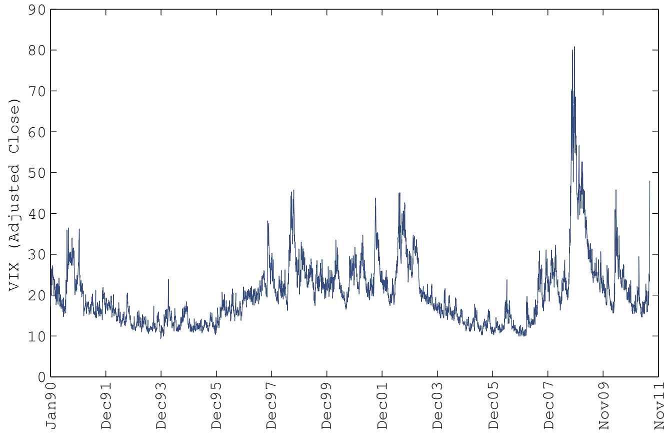

Figure 2: VIX time series. Over 20 years of history, it shows a clear mean reverting pattern to a level of around 20%.

Figure 3 shows both the collateralised[6] and uncollateralised EPE profiles for a simple spot forward and variance swap. PFE profiles show the same behaviour. The risk profile of each individual trade is shown on the left column. The netting effect is displayed in the right column by comparing added profiles versus profiles of trades which were netted first. The model was calibrated to S&P500.

The resulting netting benefit is substantial. The peak of the profile more than halves when using the spot-volatility diffusion model. An equivalent behaviour is observed for PFE profiles. This means that, without a spot-volatility diffusion model, banks calculations for CVA, exposure management and regulatory capital are notably higher than what they should be.

Model Reactivity

Another important limitation of a simple normal models is how little they react to swinging market conditions. As we are going to see, this is not only important for the three elements just mentioned above, but also for Initial Margin calculations.

Let's say that today we are living in a stressed market environment. The market points at a short-term 1 month volatility 60%, a medium-term 5-years volatility of 40% and a long-term 10-year volatility of 35%. These can be either historical or market-implied. A simple normal model will typically be calibrated to the 5-year 40% volatility point. The spot PFE profiles, both collateralised and uncollateralised for a GBM model can be

[6]Building a collateral algorithm is well beyond the purpose of this paper, so we are using 10-day risk as an indication of collateralised risk. This is equivalent to assuming an ideal CSA with daily margining, zero threshold, zero minimum transfer amount, no independent amount, no rounding, etc. The reader should note that this 10-day risk is a typical Basel's Margin Period of Risk in the context of regulatory capital calculation.

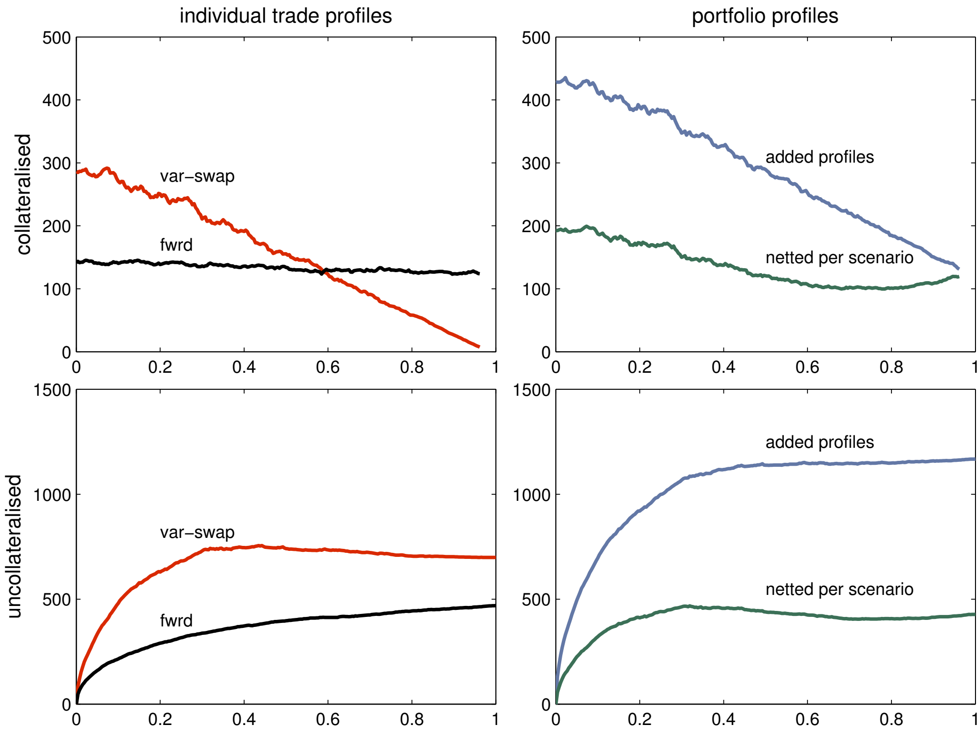

Figure 3: EPE profiles for a forward (black) and a variance swap (red), with both profiles "added" (blue) and with trades properly netted per scenario (green). This is shown both for the collateralised (top panels) and the uncollateralised (bottom panels) cases. The reader can see how the netting benefit is very important.

found in Figure 4, left panels. However, if we use a stochastic volatility model, the risk profiles are more granular in terms of volatility profile and, as such, the current high volatility can be taken into account (right panels, green lines).

Let's say that some time has passed, the source of market stress has disappeared and now we have a short-term volatility of 40%, with the same medium and long term volatilities as before. A simple normal model will not see any change in the risk profiles, but a stochastic volatility model will adapt, lowering strongly short term risk for collateralised trades (from 18 to 12 in Figure 4) and lowering long-term risks for uncollateralised profies (from 160 to around 150 in Figure 4)[7].

Then, let's say that the market went to a very quiet phase so that short term volatility was only 30%, with the same medium and long term volatility. In this case, the risk profiles will continue to decrease short-term risk under a stochastic volatility framework

[7]This happens because uncollateralised profiles are driven by the rolling 10-day risk, which will be driven by the starting volatility at the beginning of the simulation and by the mean reversion level far in time. However, in the uncollateralised world, risk is driven by the cumulated volatility up to the simulation time.

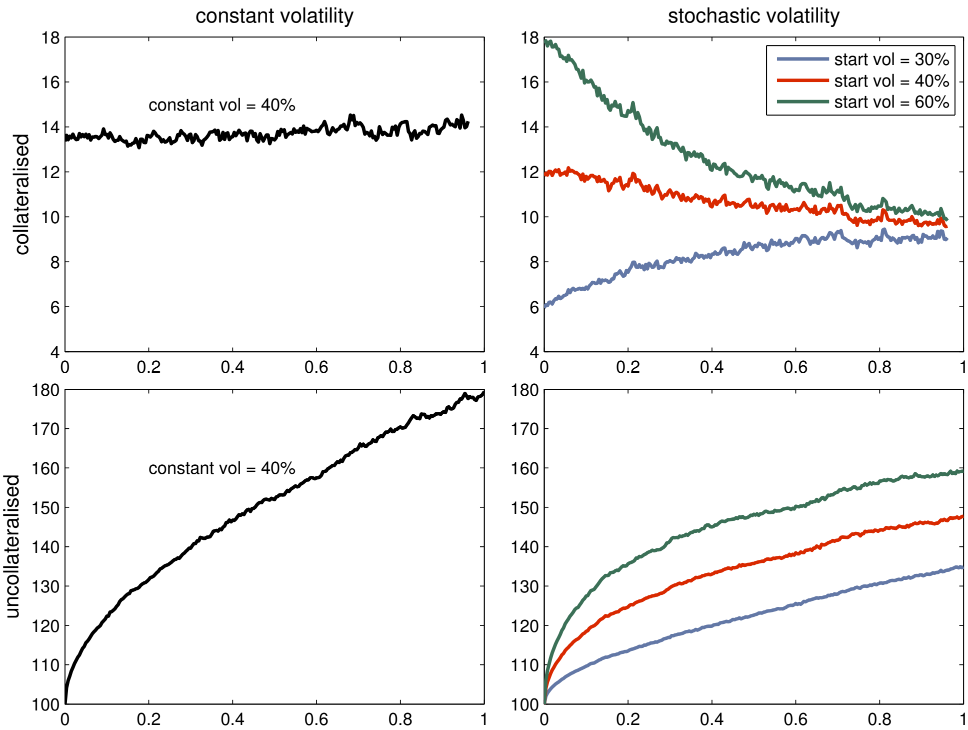

Figure 4: 95% PFE of spot, using constant volatility in a simple GBM model (left panels) and stochastic volatility in a BK spot-volatility diffusion model (right panels), displaying both the collateralised (top panels) and the uncollateralised (bottom panels) cases. The stochastic volatility simulations each use a different volatility at the start of the simulation, all other parameters being equal.

(blue lines), but will not change at all in a standard normal model.

This "reactivity" is a good feature in a model. Normal models do not have it. Stochastic volatility models do. The areas in a financial institution that are positively impacted include:

1. **CVA pricing**

   CVA is, to first order, approximately proportional to the area under the EPE profile. It can be clearly inferred from Figure 4 that that area is going to be different for different market conditions when we use a stochastic volatility model, specially for uncollateralised portfolios, reflecting market swings.

   This time let's use a interest rate swap as an example. Figure 5 shows the EPE profiles for an uncollateralised vanilla swap, modelled with a constant volatility model and a stochastic volatility model, calibrated this time to USD market implied volatilities, and to three potential market conditions: stress, normal and quiet[8].

[8]For the sake of illustration, the interest rate model was simplified to a simple 1-factor GBM process.

The bottom left panel shows the CVA for this trade under each model[9]. The difference in the CVA price delivered by each model can be clearly appreciated: the stochastic volatility model is more precise to measure CVA as it contains information about the "current" state of the market; as a result, the CVA price adapts quickly to stress or quiet conditions.

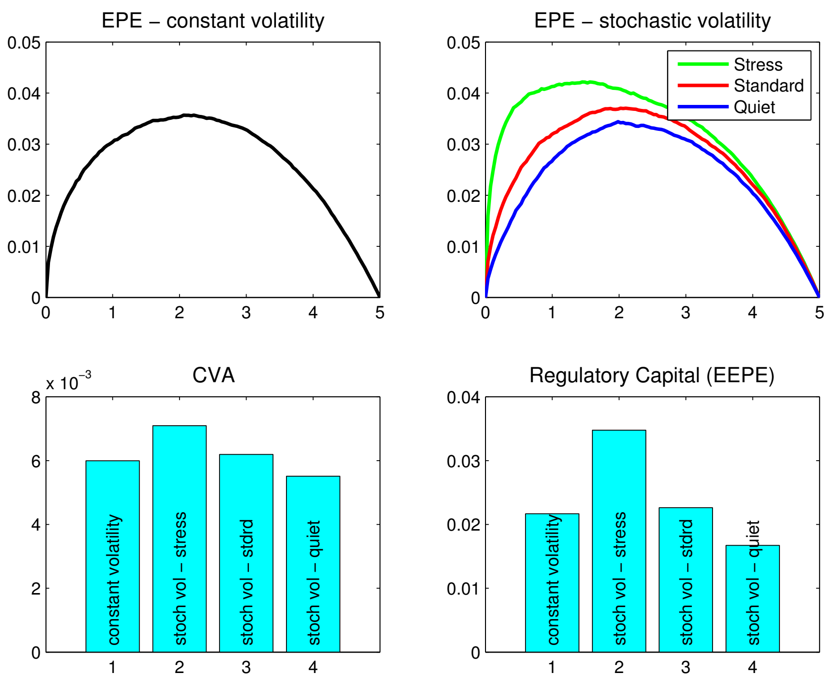

Figure 5: EPE uncollateralised profiles for an IR swap, modelled with a constant volatility model (top left), and a stochastic volatility model (top right). In the case of stochastic volatility, the model was calibrated to typical stress, normal and quiet market conditions. The bottom left panel shows the CVA for each of the four cases, and the bottom right shows Regulatory Capital (EEPE).

2. **Regulatory Capital**

Regulatory capital calculations are based in the EEPE[10] of each netting set which is, basically, the average of the non-decreasing EPE profile during the first year. A system calculating EEPE under a stochastic volatility model should in general better reflect the true 1-year default economic risk than under a constant volatility

There are other models more appropriate for interest rates, that include mean reversion for example, but the authors do not want to divert the attention from the main topic of this paper (stochastic volatility), so this model was chosen for illustrative purposes. These results can be extrapolated to more sophisticated interest rates models.

[9] CVA was calculated considering both the asset and liability side of it (asset side also known as one-way CVA, and liability side as DVA) with a flat spread of 500 bps for the counterparty and 50 bps for the bank.

[10] The regulatory Effective Expected Positive Exposure[1].

model given its better market reactivity. Figure 5 shows how the EEPE differs between both models; the reader can see how those changes can be substantial.

3. **Initial Margin Calculation**

Typically, initial margin calculations for collateralised trades are based in the maximum of the PFE profile at, for example, 99% confidence level. We can use the same swap of the previous example to illustrate the strong impact the volatility modelling has in this matter. Figure 6 shows the PFE profiles at 99% confidence for a constant volatility and a stochastic volatility model with the three calibrations mentioned before, as well as the initial margin (IM) as given by each of the models. As expected, the IM of the swap incepted in stressed market conditions and modelled under stochastic volatility is noticeable higher than that given by constant volatility. Interestingly, when the market conditions are quiet, the stochastic volatility model considers the possibility that the volatility may increase in the future and, hence, the IM is not noticeably lower compared to a constant volatility framework.

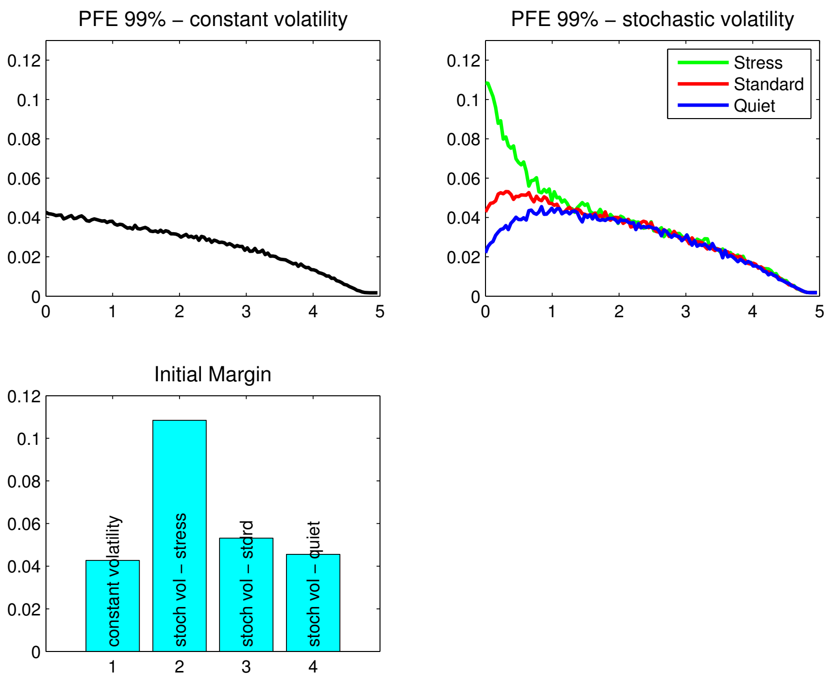

Figure 6: PFE collateralised profiles at 99% confidence for an IR swap, modelled with a constant volatility model (top left), and a stochastic volatility model (top right). In the case of stochastic volatility, the model was calibrated to typical stress, normal and quiet market conditions. The bottom left panel shows the Initial Margin calculated from each of the four cases.

This example illustrates quite clearly how the calculation of IM is more accurate with a stochastic volatility model than with constant volatility. This IM calculations are becoming increasingly important given the current trend to move to collateralised agreements and to use central clearing houses.

Better Model for Tail Dynamics

As previously mentioned, market data analysis indicates that spot time series tend to show skew and positive excess kurtosis. This is not captured at all by a simple normal processes. This is particularly important for exposure management as, here, we tend to be very sensitive the the tails of the spot distribution. Also, a model that captures well the tails should, in general, back-test better.

As said, it is well known in the academic literature that stochastic volatility models create skew and kurtosis in the spot distribution. The aim of this paper is not to do a full detailed check of this feature, but some comparison with real market data is healthy at this stage, at least to cross-check that the outcome of the mathematics make sense with reality[11].

In order to do this we calibrated our stochastic volatility model to 10 years of VIX time series, we run then a simulation of the spot and, then, we compared the 10-day standard deviation, skew and kurtosis delivered by the model to that observed in the market. Also, we compared to the results obtained from a GBM process, quite a market standard in the industry for equity modelling in counterparty risk. The following table shows the results:

10-day moves	S&P500	Stochastic Volatility	GBM
Standard Deviation	0.039	0.042	0.04
Skew	-0.89	-0.60	0
Kurtosis	9.1	5.5	3

The reader can see how a stochastic volatility model matches much better real behaviour than a GBM model. This test has been done only for one time series (S&P500), so it should not be extrapolated naively and say that a stochastic volatility model will always be better in this context than a GBM, but it is a healthy and encouraging cross-check.

In any case, it appears that stochastic volatility models may be better than normal processes specially for, for example, exotic trades, where risk can be highly sensitive to the tails of the spot distribution.

Calibration and Some Practicalities

Typically, counterparty risk models used for the purpose of exposure management and regulatory capital are historically calibrated, and when for the purpose of CVA pricing

[11] For example, it would be absurd to find that the calibrating parameters deliver a skew with opposite sign of that observed, or a kurtosis 10 time higher that that seen in reality.

are market-implied calibrated. These stochastic volatility models can be calibrated to both schemes.

If for CVA, we can calibrate, for example, "today's" volatility to the 1-month at-the-money option implied volatility, and θ and μ can be calibrated as a best fit of the expected future volatility to the term structure of the implied volatilities. The vol-of-vol term η will typically need some sort of historical analysis (unless there is a liquid options-of-options market); for example, we can calibrate η to the standard deviation of changes in the short-term point of the volatility surface.

If historical calibration is needed, then we can calibrate "today's" volatility to the last 1-month standard deviation of daily spot moves and then other parameter like θ and μ using the expected future value of the volatility with a fit it to a number of points of the long-run historical volatility (e.g., 1, 5, 10 years). The vol-of-vol is, arguably, more tricky the calibrate; we can use, for example, a long time series of 1-month rolling volatility for it, and calibrate it to the standard deviation of changes, considering the relevant correcting factor if the data contains overlapping information[3].

An interesting mixed version of both schemes can also be used. We can pick the most representative point of a options volatility surface (e.g., 1-year at-the-money), calculate a long time series of it, and calibrate all volatility parameters to it using the standard least-squares or maximum likelihood methods; that is, we are using a history of market-implied data. In fact, in reality, this type of mixed historical-implied methodologies will need to be used often as a result of lack of available data.

There is no given calibration methodology that is, by definition, better than the other ones. Which methodology to use will depend in the final use of the model, internal policies and, most importantly, in data availability.

In addition to calibration, a practitioner designing a stochastic volatility model for counterparty risk should consider a number of practicalities, like the sensitivities that the portfolio under consideration has an how to apply these models to different asset classes. The examples of this paper have been focused in 1-factor models. We have calibrated our illustrative model to Equities (S&P500) and to interest rates (USD). These stochastic volatility models can be applied to any asset class, as all of them suffer in some way or another from the limitations highlighted here when modelled with constant volatility applied to a Brownian motion. However, multi-factor stochastic volatility models can become quite convoluted in practice. As such, the researcher will have to weight out the benefits and costs of stochastic volatility. In order to do this, the researcher has to understand what are the main risk factors that the portfolios under consideration are sensitive to (e.g., is my portfolio sensitive to rotations in the yield curve?) and find, on this way, the optimum balance between risk measurement quality, model complexity and implementation practicalities.

Summary, Conclusions and Critique

As shown in this paper, migration from constant volatility to more advanced volatility-spot stochastic models can deliver important benefits for counterparty risk management, capital calculation, and CVA pricing.

In contrast to typical normal or log-normal models, where the volatility is constant, we showed how spot-volatility diffusion models provide strong netting benefits between spot and volatility risk-driven trades. Then we showed how these models can differentiate between short-term and long-term risk, providing better CVA, capital and initial margin measurements. Finally, we showed how these models introduce skew and kurtosis, a desirable feature.

Together with the main strengths outlined, these frameworks have a number of potential weaknesses which must be kept in mind. As seen, the kurtosis it delivers seems to be still too low[12], so this model may not be appropriate yet to measure extreme tail risk. Also, for a given time horizon and depending in the model, the probability distribution of the spot under a spot-volatility diffusion process can depend on the number of time steps taken by the simulation. However, we found that that dependence was small in the illustrative model we used. Finally, more sophisticated correlation structures between the volatility and the spot could be investigated, as we only implemented a locally gaussian copula dependency structure.

Acknowledgements

The authors would like to thank Piero Del Boca, Keith Garbutt, Seung Yang and both referees for interesting remarks on this paper.

References

[1] *International convergence of capital measurement and capital standards*, tech. rep., Bank for International Settlements, 2006.

[2] L. B. G. Andersen, *Efficient simulation of the heston stochastic volatility model.* Working Paper, 2007.

[3] P. Bod, D. Blitz, P. H. Franses, and R. Kluitman, *An unbiassed variance estimator for overlapping returns*, Applied Financial Economics, 12 (2001), pp. 155–158.

[4] P. Carr and D. B. Madan, *Option valuation using the fast fourier transform*, Journal of Computational Finance, 2 (1999), pp. 61–73.

[5] J.-P. Fouque, G. Papanicolaou, and K. R. Sircar, *Mean-reverting stochastic volatility.* Working Paper, 1999.

[12] At least in the particular case we studied.

[6] L. Gauthier and D. Possamay, *Efficient simulation of the double heston model*. Working Paper, 2011.

[7] L. A. Grzelak and C. W. Oosterlee, *On the heston model with stochastic interest rates*, Journal of Financial Mathematics, 2 (2011), pp. 255–286.

[8] S. Mikhailov and U. Ngel, *Heston's stochastic volatility model implementation, calibration and some extensions*, Willmott Magazine, (2005).

[9] E. W. Weisstein, *Student's t-distribution*. From MathWorld - A Wolfram Web Resource. http://mathworld.wolfram.com/Studentst-Distribution.html.

[10] P. Wilmott, *Paul Wilmott on Quantitative Finance*, Wiley, second ed., 2006.