MOODY'S ANALYTICS

B&H RESEARCH
SEPTEMBER 2013

Steven Morrison, PhD
Laura Tadrowski, PhD
Moody's Analytics Research

Contact Us

Americas
+1.212.553.1658
clientservices@moodys.com

Europe
+44.20.7772.5454
clientservices.emea@moodys.com

Asia (Excluding Japan)
+85 2 2916 1121
clientservices.asia@moodys.com

Japan
+81 3 5408 4100
clientservices.japan@moodys.com

Guarantees and Target Volatility Funds

Overview

Over the last five years, target volatility funds – where an asset mix is dynamically re-balanced with the aim of maintaining a stable level of portfolio volatility through time - have emerged as an increasingly popular asset class. In both North America and Europe, major Variable Annuity (VA) writers have included or are considering the inclusion of target volatility funds as underlying investment funds for their VA products. The dynamic risk management that is embedded in these funds naturally controls guarantee costs, whilst still offering the policyholder upside exposure to risky assets. At a time when policyholders appear reluctant to pay (market-consistent) guarantee costs, these funds are therefore naturally attractive to VA product designers.

However, assessing the guarantee costs that are produced when invested in target volatility funds is not straightforward. In particular, the guarantee cost assessment depends critically on modeling assumptions about how well target volatility funds can do what they say on the tin. Specifically, can these funds always deliver stable volatility levels, or are there sources of equity market volatility which cannot be controlled by the funds' dynamic re-balancing? And if so, what are the consequences for guarantee costs?

In this paper we consider the valuation of guarantees written on target volatility funds, and the sensitivity of these values to the choice of equity model and rebalancing frequency.

Contents

1. Introduction

The high levels of equity implied volatility experienced during the recent global financial crisis, combined with relatively low levels of risk-free interest-rates, significantly increased the market-consistent cost of guarantees embedded in Variable Annuity (VA) products. As a result, VA providers have sought methods that retain the guarantees that policyholders find so attractive, while reducing their costs to affordable levels. One such method, that has gained popularity in recent years, is the use of target volatility funds. The basic idea of a target volatility fund is that the investment fund underlying the VA product is dynamically rebalanced so as to achieve a certain target level, thereby reducing the cost of guarantees written on this fund.

Product providers commonly use stochastic models to price guarantees and to calculate price sensitivities for the purpose of hedging. The aim of this paper is to investigate the impact of some of the key modeling decisions on the cost of guarantees. For illustration, we consider the pricing of vanilla put options written on an underlying target volatility fund. The rest of the note proceeds as follows:

- Section 2 describes the dynamic rebalancing rules adopted in the analysis.
- Section 3 analyses how guarantee costs change as we use different models for equity returns, including models exhibiting stochastic volatility and jumps.
- Section 4 analyses how guarantee costs change as we change the frequency at which the target volatility fund is dynamically rebalanced.
- Section 5 concludes.

2. Target volatility fund dynamics

A target volatility fund is a portfolio of a risky assets (typically equities) and risk-free assets, with the proportion invested in each dynamically rebalanced so that the volatility of the fund is equal to, or close to, some target volatility.

In this note, we consider a typical dynamic rebalancing strategy, where the weight in equity at rebalancing time t is set to:

$$w_t^{equity} = min\left(\frac{\sigma_{target}}{\hat{\sigma}_t^{equity}}, 100\%\right)$$

where σ_{target} is the target volatility and $\hat{\sigma}_t^{equity}$ is an estimate of the volatility of equity returns applying between time t and the next rebalancing time. We constrain the maximum weight in equity to be 100% (i.e. if estimated equity volatility is too low, we don't leverage in order to match the target volatility). With this strategy, the volatility of the fund over each rebalancing period is exactly equal to the target volatility, assuming we can perfectly estimate equity volatility and assuming that this is always greater than the target.

In practice, we cannot estimate future equity volatility exactly. Volatility is typically estimated using historical equity returns, with Exponentially Weighted Moving Average (EWMA) estimators being widely used. The EWMA estimate of the volatility of equity log-returns applying between time t and $t + \Delta t$ is:

$$(\hat{\sigma}_t^{equity})^2 = \lambda(\hat{\sigma}_{t-\Delta t}^{equity})^2 + (1-\lambda)\frac{1}{\Delta t}\left(ln\left(\frac{S_t}{S_{t-\Delta t}}\right)\right)^2$$

where S_t is the equity index at time t. In the examples below we will choose $\Delta t = 1$ business day (with 252 business days per year). This estimate puts progressively less weight on older data, with the weight decaying exponentially. λ controls the rate of exponential decay, with higher values of λ corresponding to slower rates of decay, and the mean age of the data used being $\frac{\Delta t}{1-\lambda}$ years. In the examples below we choose $\lambda = 0.99$, corresponding to a mean age of 0.4 years. This parameterization is representative of the assumptions typically used by target vol funds in their dynamic re-balancing algortihms.

Note that, to the extent that we cannot estimate volatility exactly, the volatility of the fund will not exactly equal the target volatility – if the EWMA estimator happens to overestimate equity volatility we will allocate too little weight in equity and so the fund volatility will fall short of the target, while if it underestimates volatility we will allocate too much weight in equity and so the fund volatility will exceed the target. We will see below that this estimation error can have a significant impact on guarantee costs, depending on the choice of equity model.

3. The effect of equity model choice

Firstly, we consider the sensitivity of option prices to different choices of model for the equity fund. We assume that interest rates are deterministic, and consider three different models for returns in excess of the risk-free short-rate:

1. A constant volatility (Black-Scholes) model.
2. A stochastic volatility (Heston) model.
3. A model incorporating both stochastic volatility and jumps (Stochastic Volatility Jump Diffusion or SVJD).

For technical details of the Heston and SVJD models, see (Calabrese, Gawlikowicz and Lord 2013).

In all cases, we assume that the fund is rebalanced each business day, equity volatility is estimated using an EWMA estimator with $\lambda = 0.99$, and the target volatility is 10%. Each model has been calibrated to implied volatilities on EuroStoxx 50 index options at

end-December 2012[1]. Similar analysis was recently carried out by (Jaschke 2013), albeit using different model specifications for stochastic volatility and jumps.

Figure 1 shows estimated implied volatilities for 2, 5 and 10 year options at a range of strikes using the constant volatility model. Implied volatilities are shown for put options written on the Target Volatility fund, and on the underlying equity index for comparison. Put option prices were estimated using 5,000 scenarios[2], generated using the Barrie & Hibbert Economic Scenario Generator.

Figure 1: Implied volatilities – Constant volatility model (Black-Scholes)

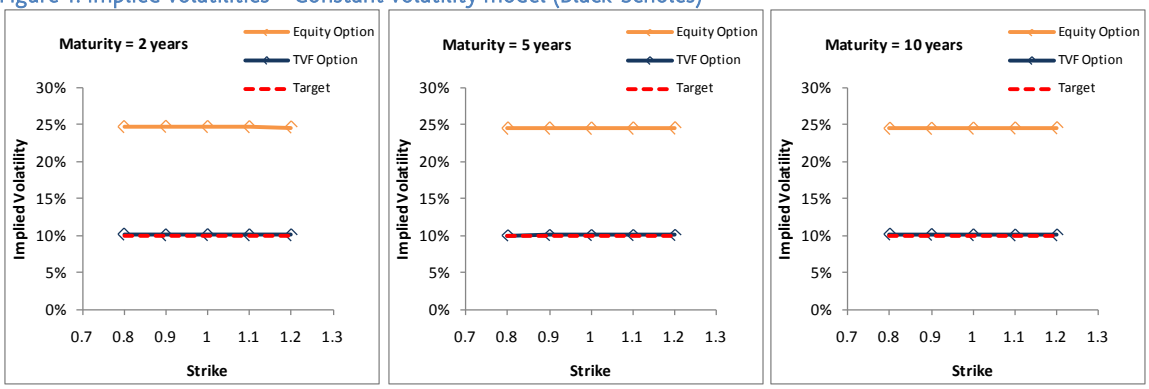

In the case of the constant volatility model, the implied volatility on the target volatility fund is approximately equal to the target volatility at all maturities and strikes. Though this result may be intuitive and expected, it will be helpful in comparing with later models to explain using a sample scenario.

Figure 2 shows a sample scenario for the equity index and its volatility and the resulting fund value and its volatility. We observe that the estimated (EWMA) volatility remains stable and close to the 'true' model volatility (here assumed to be 24.5%, consistent with long maturity EuroStoxx 50 market implied volatilities at end-December 2012). In this case, the EWMA estimator tends to do a good job of estimating the true model volatility. As a result, the volatility[3] of the fund is also reasonably stable and close to the target volatility. Since the fund is rebalanced frequently, and its volatility is approximately constant and equal to the target volatility, the fund dynamics are well approximated by the Black-Scholes model and hence the implied volatility is approximately equal to the target volatility at all maturities and strikes.

[1] Constant volatility and SVJD models were calibrated using Moody's Analytics standard calibrations of these models (Barrie & Hibbert 2013), while the Heston model was calibrated to be consistent with the same option implied volatility data at the 2, 5 and 10 year maturities considered here.

[2] Standard errors are not shown, typically being small compared to the symbols on the charts here.

[3] When we refer to volatility here, we mean the 'true' model volatility of the fund.

Figure 2: Example scenario – Constant volatility model (Black-Scholes)

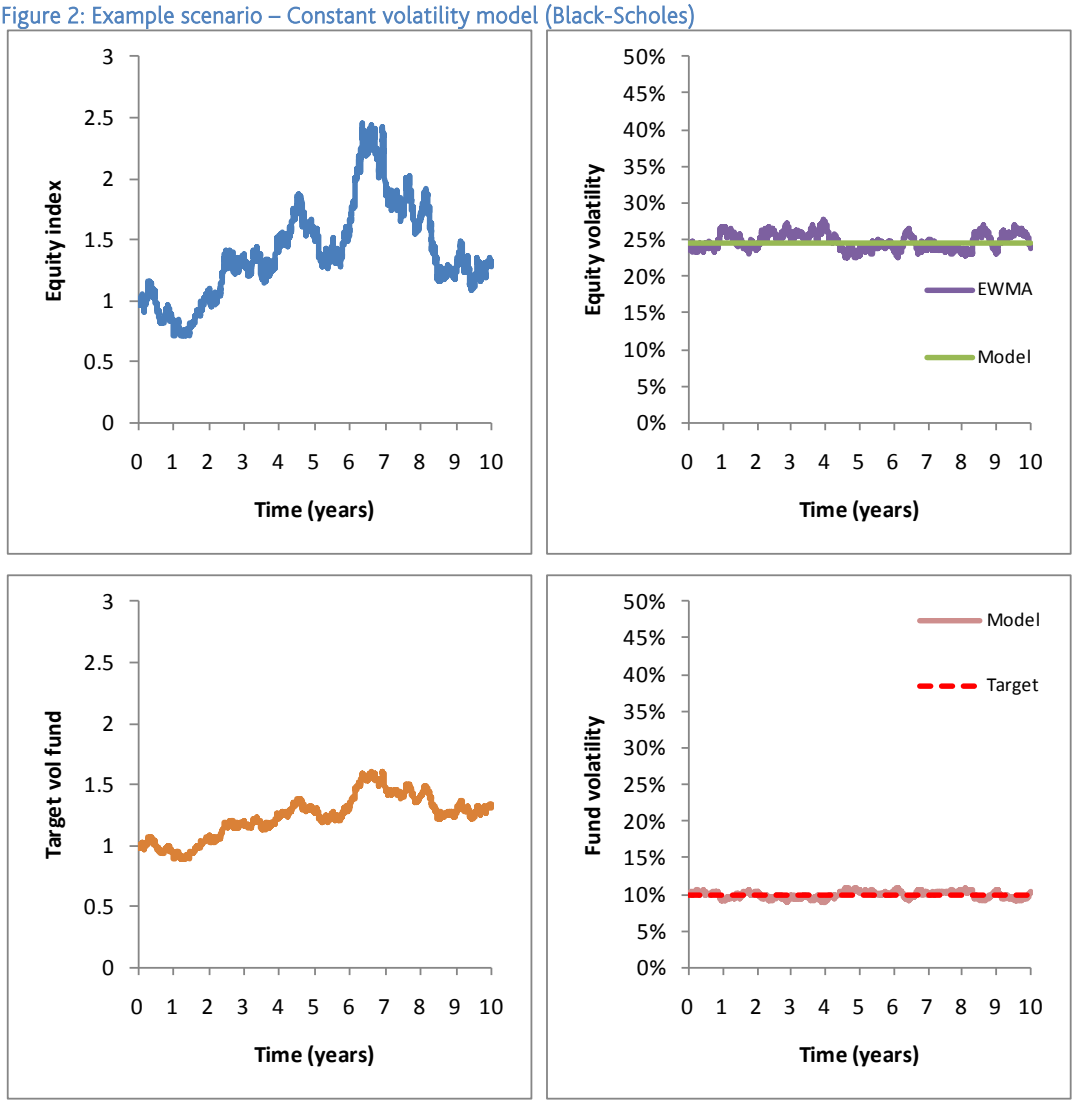

We now move on to the more complex stochastic volatility model (but not yet with jumps). The put option implied volatilities for this model are shown below in Figure 3.

Figure 3: Implied volatilities – Stochastic volatility model (Heston)

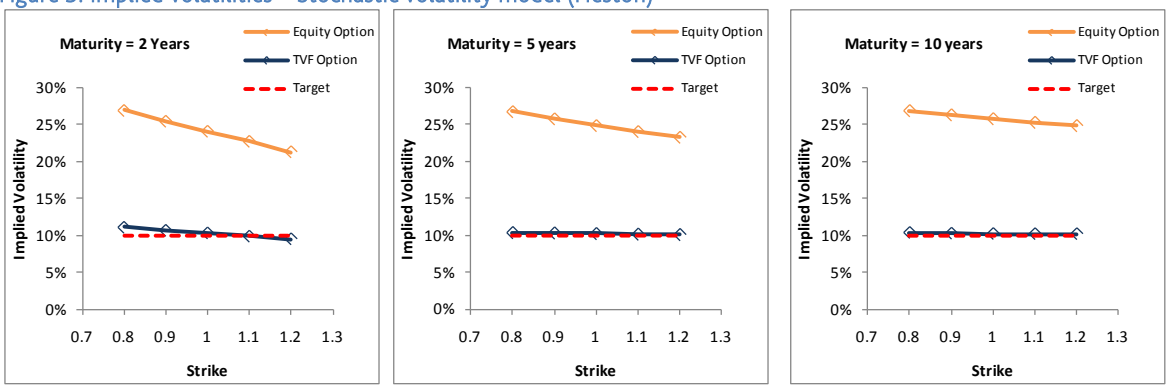

The implied volatility on the target volatility fund is still approximately equal to the target volatility at all maturities and strikes considered, with only a slight skew observed at the 2 year maturity. For the stochastic volatility model, Figure 4 shows a sample scenario for the equity index and its volatility and the resulting fund value and its volatility.

Figure 4: Example scenario – Stochastic volatility model (Heston)

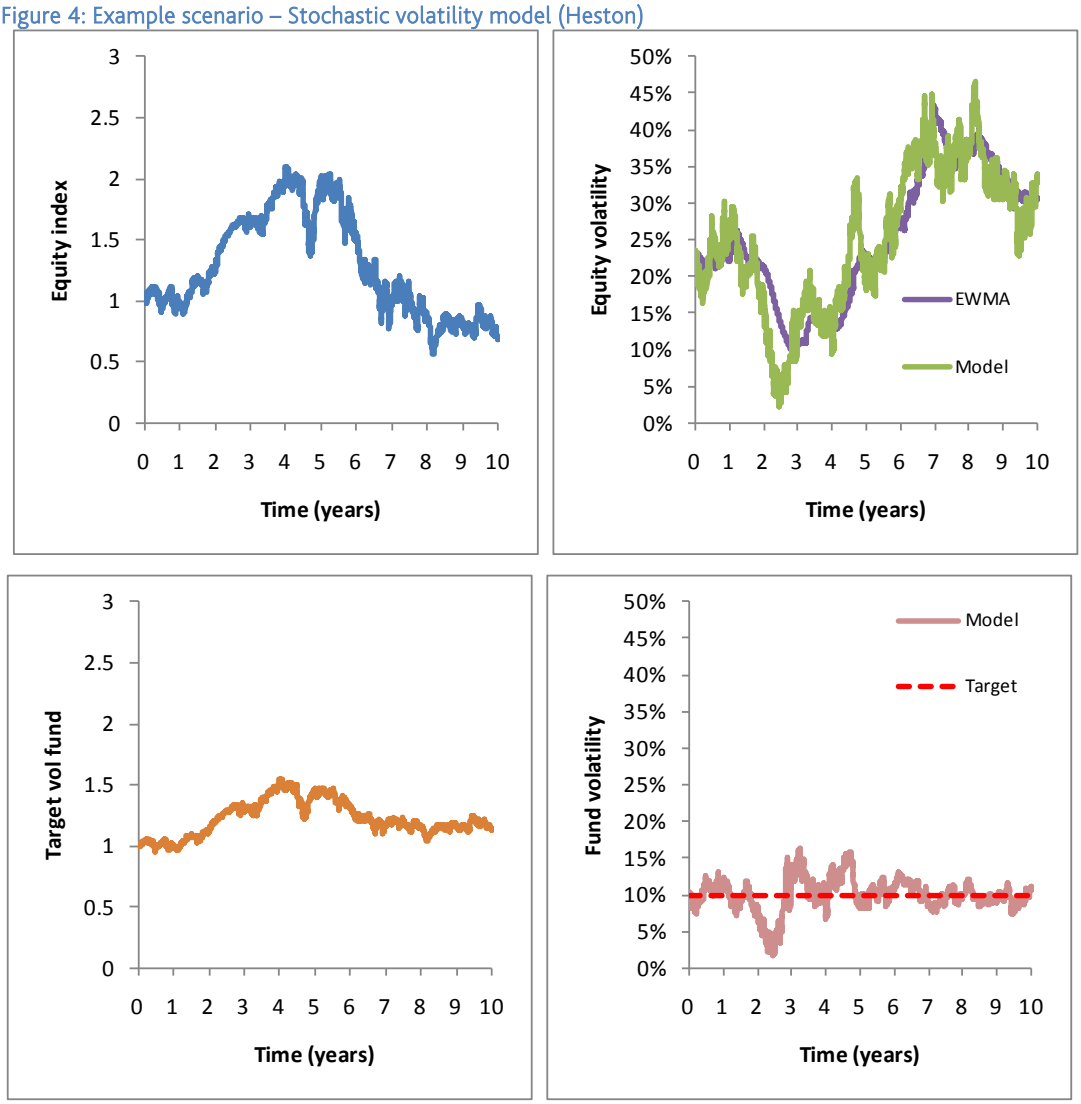

Again, we observe that the EWMA tends to do a good job of estimating the true model volatility, but here the equity model volatility is stochastic, varying in the range 2.3% to 46.6% in this particular scenario. We also observe that the EWMA estimator tends to 'lag' behind the true model volatility, due to the fact that it is always estimated using 'old' data (albeit with an exponentially decaying weighting). The size of this lag depends on the chosen λ parameter and we would expect smaller λ (and hence more weight on more recent data) to reduce the lag (at the expense more 'noise' in the estimate)[4].

The resulting fund volatility is again approximately equal to the target volatility on average, and as a result implied volatilities are approximately equal to the target volatility. However, the estimate is less stable (i.e. average deviations from the target volatility are larger) than in the constant volatility model. This is largely due to the lag in the EWMA estimate and the fact that the underlying 'true' model volatility is stochastic. For example, in this scenario we see a sharp decrease in modeled volatility around year 2. At this time, the lag in the EWMA estimator results in an overestimation of model equity volatility (since the EWMA estimator puts significant weight on returns produced during the earlier, higher volatility, period), we underweight in equity, hence the volatility of the fund is lower than the target. Nonetheless there is far less variation in the volatility of the fund than in the volatility of the underlying equity index. As a result, while there is significant skew in market implied volatilities on the equity index, the implied volatilities on the fund are far flatter as a function on option strike, and by 5 year maturity the skew is almost completely removed.

We now move on to the SVJD model. Implied volatilities are shown below in Figure 5.

[4] In theory, we could set the equity fund's volatility forecast to be equal to the 'true' model volatility, so as to remove any estimation noise. In practice, the true model volatility is unknown and must be estimated, either using historical estimators or implied volatility, so that some estimation error is inevitable.

Figure 5: Implied volatilities – Stochastic Volatility Jump Diffusion (SVJD) model

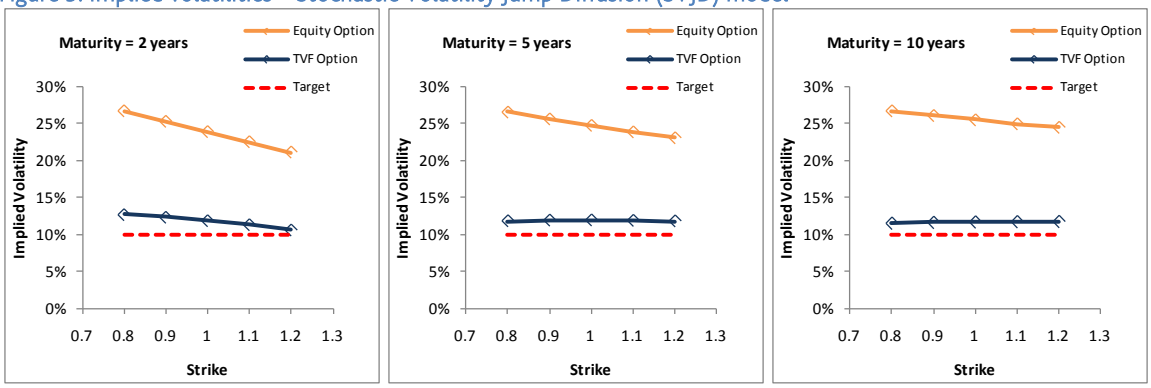

The option prices for the target vol fund are higher than that implied by the target volatility at all maturities and strikes considered, and there is increased skew in implied volatilities at the 2 year maturity (though again, this appears to flatten off by year 5).

Figure 6 shows an example scenario for the SVJD model.

Figure 2: Example scenario – Stochastic Volatility Jump Diffusion (SVJD) model

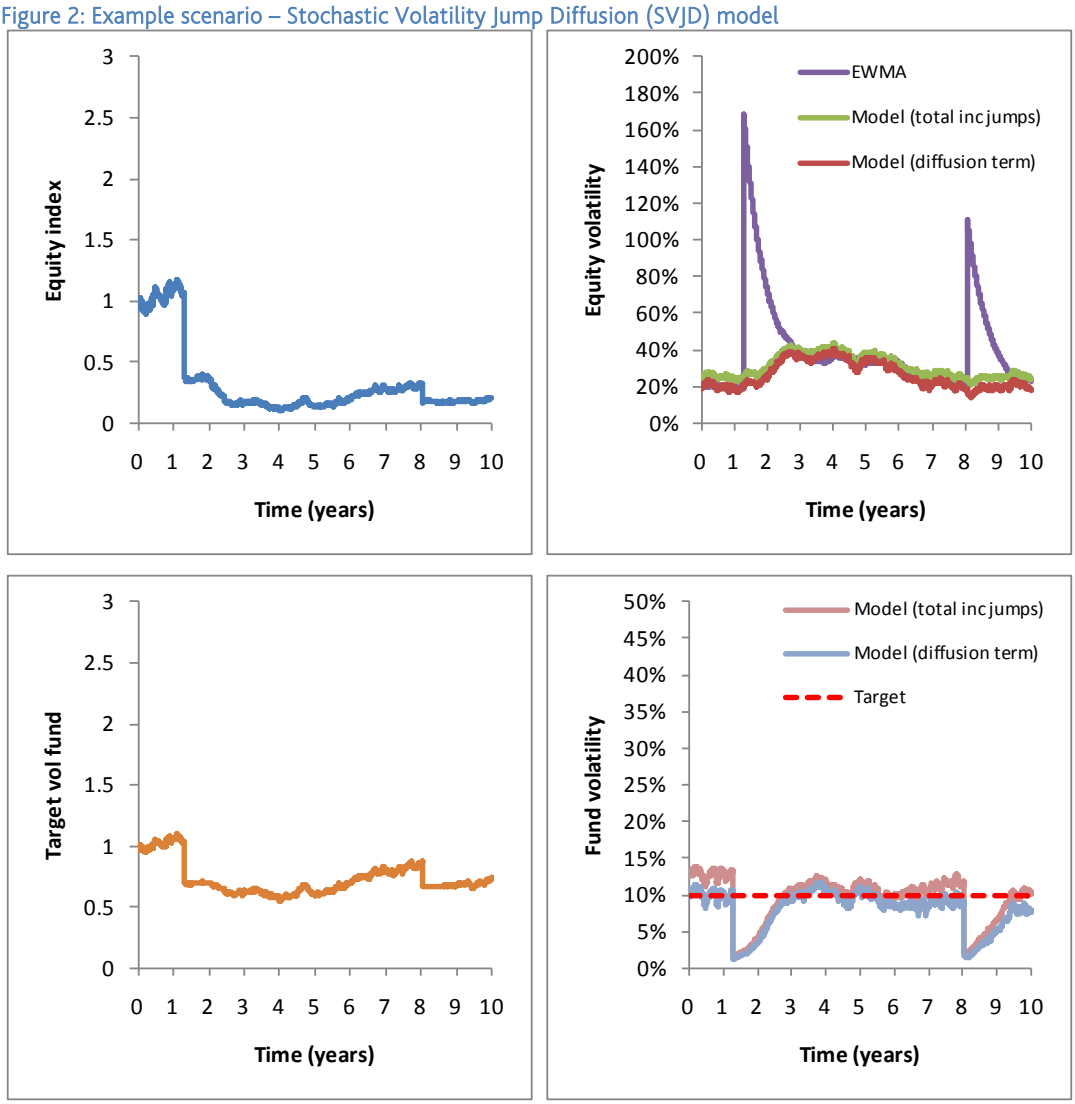

The example scenario path in Figure 6 contains two equity jumps (at around 1.3 and 8 years). The charts again compare the EWMA estimate of equity volatility to the model volatility, now defined in two different ways:

» 'Model (diffusion term)' is the instantaneous volatility of returns as implied by the 'diffusion' term in the model i.e. the stochastic volatility. This represents the instantaneous volatility of returns in the absence of jumps.

» 'Model (total inc jumps)' is the total instantaneous volatility including the effects of both stochastic volatility and jump terms, with jumps contributing additional volatility.

During periods in which there are no jumps, the EWMA estimate again does a reasonable job of estimating the 'true' model volatility (the stochastic volatility). However, immediately following an equity jump, there is a subsequent sharp increase in the EWMA estimate of equity volatility as the (typically large) jump return suddenly appears in the EWMA estimator. As a result the target weight in equity, and hence volatility of the resulting fund, falls sharply, and for a period following the jump is significantly lower than the target volatility.

Due to these periods of relatively low fund volatility *following jumps* we might expect implied volatility to be lower than the target volatility, but in Figure 5 we observe that the implied volatility is higher than the target volatility. At-the-money implied volatilities on the fund are estimated as 11.9% (at maturities of 2 and 5 years) and 11.7% (at maturity of 10 years). However, *at the time of the jump*, our EWMA estimator of equity volatility is too low – this estimate is based on a set of historical returns that don't contain any jumps and so only estimates the volatility of the diffusion term in the model, rather than the total volatility including jump returns. By not accounting for jumps, we underestimate equity volatility before the jump occurs, with the result that we are overweight in equity at that time. The net effect is that we tend to overweight equity and the resulting implied volatility on the fund is larger than target volatility.

Note also that returns on the target volatility fund also contain jumps. Indeed, no matter how we rebalance the fund, as long as we are exposed to equity so the return on the fund contains jumps. This, combined with the fact that the volatility of the fund is stochastic (despite our attempts to control it) gives rise to skew in the distribution of returns and hence in option implied volatilities – at the 2 year maturity implied volatilities on the fund are estimated as 12.7% (strike =0.8) and 10.6% (strike = 1.2) – though as in the pure stochastic volatility model this skew appears to flatten as we increase the maturity.

These results confirm those observed in (Jaschke 2013), albeit using different models for stochastic volatility and jumps, and lead to the conclusion that a model with jumps suggests a significant larger market-consistent cost (and larger skew at short maturities) for target volatility fund guarantees than implied by a model without jumps, even if these are calibrated to the same underlying market equity option implied volatilities. The presence of jumps significantly complicates the estimation of equity volatility with the result that we tend to overweight in equity using standard dynamic rebalancing strategies.

Figure 7 compares (at-the-money) put option prices under the SVJD model with those produced by the equivalent stochastic volatility model without jumps. The inclusion of jumps results in option price increases of 18% (2-year option), 24% (5-year option) and 28% (10-year option).

Figure 3: Prices of at-the-money put options on a Target Volatility Fund, under models with and without jumps

Model	2 Year	5 Year	10 Year
No Jumps	0.0458	0.0556	0.0538
With Jumps	0.0542	0.0689	0.0688
% Increase	+18%	+24%	+28%

4. The effect of rebalancing frequency

So far, we have assumed that the target volatility fund is rebalanced each business day, which is typical of many funds in practice. In this section, we explore the sensitivity of results to rebalancing frequency.

Here we assume that equity excess returns follow the SVJD model, calibrated to implied volatilities on EuroStoxx 50 index options at end-December 2012. In all cases, we estimate (daily) equity volatility using an EWMA estimator with $\lambda = 0.99$, and the target volatility is 10%. However, we now vary the frequency at which we rebalance the fund so as to achieve this target.

Figure 8 shows estimated fund implied volatilities for 2, 5 and 10 year options at a range of strikes. As before, put option prices were estimated using 5,000 scenarios, generated using the Barrie & Hibbert Economic Scenario Generator. Here we compare daily, weekly, monthly and quarterly rebalancing frequencies.

Figure 8: Implied volatilities at different rebalancing frequencies (SVJD model)

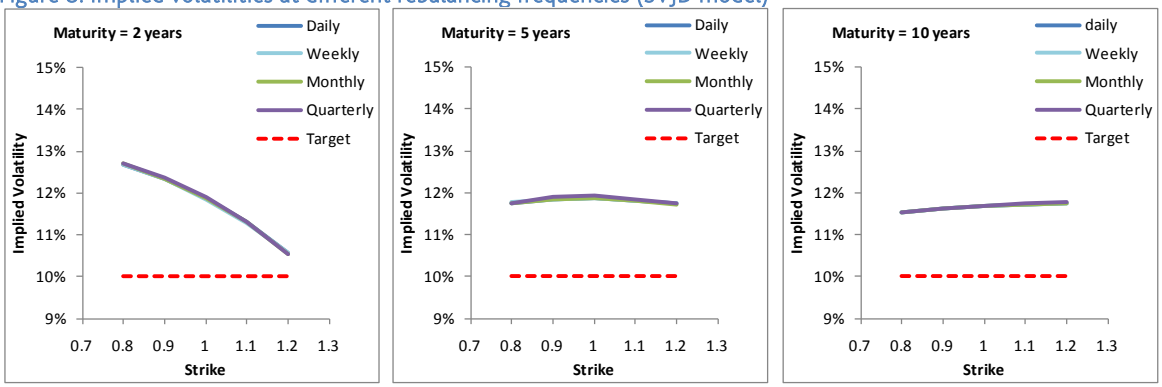

We observe that implied volatilities are relatively insensitive to rebalancing frequency over the range of frequencies considered here. To explain this result, Figure 9 shows the same sample scenario before, compared over the different choices of rebalancing frequency. The upper two charts show equity weight and level of the fund over the full 10 year projection, the middle two charts show the same data over the first 2 years only, and the lower 2 charts show the same data over the first two quarters only.

Looking at the weights, we see that these agree exactly at common rebalancing dates, by construction. However, between rebalancing dates the equity weights 'drift' according to realized equity returns. The extent to which rebalanced weights and drifting weights differ depends on realized equity volatilities and returns (and the rebalancing algorithm), but in the example scenario here we observe relatively small differences. The largest differences occur after equity jumps, but these are rare and any differences are rebalanced at the next rebalancing date. On the scale of the maturity of the options considered here, the overall differences in weights (and hence fund volatilities) are small on average and we find that the resulting fund values are relatively insensitive to rebalancing frequency. In recognition of this, some target volatility funds allow the weight to naturally drift, only rebalancing if it drifts sufficiently far from the target weight.

Figure 9: Example scenario at different rebalancing frequencies

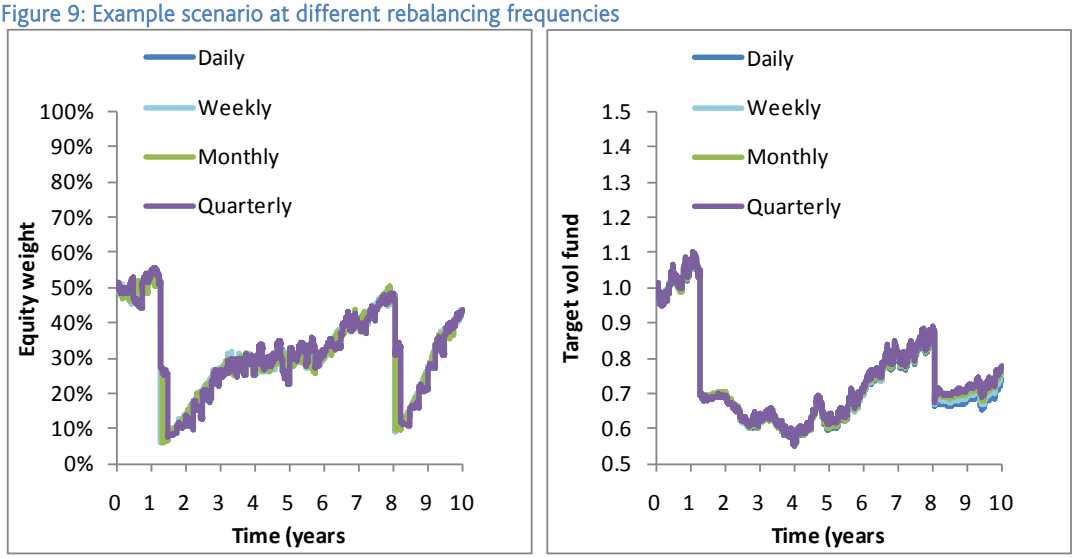

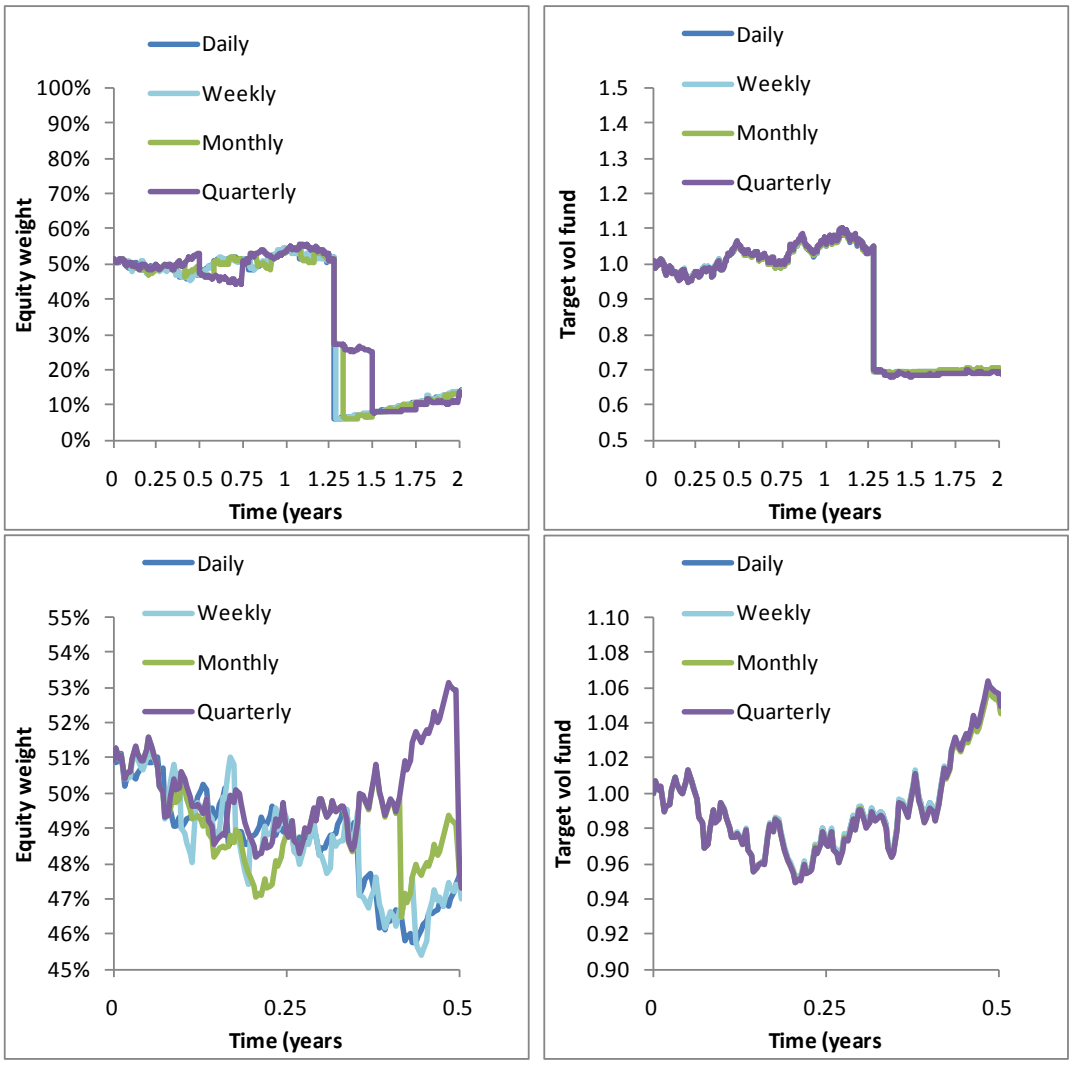

5. Conclusions

In this note we have considered the valuation of guarantees written on target volatility funds, and the sensitivity of these values to choice of equity model and rebalancing frequency.

We observe that under an assumption that equity volatility is stochastic (but in the absence of equity jumps), the target volatility mechanism works almost 'perfectly', in the sense that the resulting market-consistent guarantee costs' implied volatilities are close to the target volatility at all maturities and strikes considered. In contrast, the existence of jumps results in systematic underestimation of equity volatility and a resulting overweighting in equity, at the times when equity returns are largest (i.e. during a jump). This results in quite significant increases in guarantee costs relative to those produced by a model without jumps, even if these are calibrated to the same underlying market equity option implied volatilities – in our examples, *guarantee costs increased by 18%-28% when jumps were incorporated into the equity model. Given that jumps are a recognized feature of real short-term equity returns, this implies that it is very important to include them in the modeling of target volatility funds when assessing the cost of guarantees written on such funds.*

Furthermore, we observe that for the particular rebalancing strategy considered here, the level of the target volatility fund, and hence cost of guarantees written on it, is relatively insensitive to the assumed rebalancing frequency. This may be an important consideration for product designers, who may be able to achieve much of the benefit of the target volatility fund with less frequent rebalancing and hence lower costs, and modelers, who may be able to approximate the dynamics of target volatility fund using less frequent rebalancing and hence lower computational requirements.

References

Barrie & Hibbert. "Market consistent asset model calibrations: EUR Govt Bonds End-Dec 2012." 2013.

Calabrese, Gioel, Witold Gawlikowicz, and Natacha Lord. *Stochastic Volatility Jump Diffusion: Model Definition.* Barrie & Hibbert, 2013.

Jaschke, Stefan. "Pricing and hedging derivatives on actively managed funds." *Investment-based Insurance Guarantees.* Dublin, 2013.