



Munich Personal RePEc Archive

A stylised model for extreme shocks: four moments of the apocalypse

Brace, Allan and Lauer, Mark and Rado, Milo

National Australia Bank, National Australia Bank, National
Australia Bank

15 August 2007

Online at <https://mpra.ub.uni-muenchen.de/124773/>
MPRA Paper No. 124773, posted 23 May 2025 14:27 UTC

A Stylised Model for Extreme Shocks: Four Moments of the Apocalypse

Alan Brace, Mark Lauer and Milo Rado*

Group Market Risk,
nabCapital, Australia

15th August, 2007

Abstract

We present a method for calculating the extreme tail quantiles, over arbitrary holding periods, of a continuous-time stochastic volatility model of the form proposed by Scott (1987) with correlation between the processes for volatility and price. The fat tails of this model enable a consistent, tuneable, stylised representation of non-normality in extreme moves of prices across differing markets.

Because the model is analytically intractable, four moments are derived by numeric integration and matched to a one-period version of the model, whose quantiles are then found by further numeric integration. We also present a novel Monte-Carlo simulation scheme, which we have used to confirm the accuracy of the moment-matching approximation for quantiles as extreme as one-millionth.

Two methods for calibrating the model to market data are also proposed. The model is used in production stress testing at nabCapital to define consistent real-world probabilities for extreme shocks over heterogeneous holding periods.

1 Introduction

Stress testing is widely considered to be a necessary supplement to other risk measures such as Value-at-Risk (VaR) in best-practice risk management for financial institutions. Indeed stress testing is required by financial regulators (for example, APRA, 2000, paragraph 23). Unlike VaR, however, there are no commonly accepted industry standards for determining shocks to be applied in these stress tests. Commonly stress shocks are constructed

*All Rights Reserved, National Australia Bank, 2007. The authors would like to thank Volf Frishling and David Stump for valuable contributions. Correspondence to: <mark.lauer@nab.com.au>

from hypothetical scenarios or from historical events, with no probabilities assigned to them (CFGs, 2005). Yet, without such probabilities, it is impossible to evaluate the relevance of the scenarios (as Berkowitz, 1999, has argued), or even to compare results from different scenarios to one another.

While in principle one can assign probabilities to scenarios *ex-post*, using for example whichever model is used for VaR, this does not immediately allow comparison of scenarios when probabilities for the scenarios turn out to differ. Also, commonly used VaR models have poor performance in estimating tails more extreme than those used for VaR (Alexander & Sheedy, 2007), or cannot do so at all¹. Instead, we would like to derive the shock sizes having a given fixed probability. This ensures that all scenarios are equally relevant and important, permitting direct comparison amongst scenarios and also uniform stress test limits.

Moreover, if shock sizes with sufficiently small probabilities can be derived, then stress test results can be tied directly to economic capital requirements for a desired credit rating. Starting with a desired credit rating, we can determine the corresponding acceptable default probabilities over various holding periods, and then build stress tests whose probability equals those default probabilities. Loss results arising from these stress tests will then represent the amount of capital an institution must hold in order to keep the probability of exhausting that capital below the institution's own default probability. This leads to a natural equivalence between stress test limits and economic capital.

We therefore seek a model that is consistent with standard VaR models, but is focussed on extreme tail quantiles of arbitrarily low probability. Since regulators require stress testing to be applied to all major risk types (APRA, 2000, paragraph 24) and all material portfolios, the model must be applicable to a broad range of risk sources and be simple to apply. Because a global diversified capital markets business is exposed to hundreds or even thousands of individual risk factors, it must be possible to calibrate the model without human judgement being applied to each specific risk factor.

Liquidity is also a critical issue in stress testing, with regulators requiring explicit incorporation of this aspect (APRA, 2000, paragraph 25). Extreme market moves are often accompanied by sharp falls in liquidity, and traders may be unable to close positions for extended periods of time. A natural way to capture this is to apply holding periods according to the likely worst-case delay in closing positions (this approach is also suggested in Alexander & Sheedy, 2007). Since these holding periods vary amongst different sources of risk, the model must be able to represent arbitrary holding periods consistently; preferably, it should be independent of any specific time interval.

¹For instance, the conventional method of historical simulation cannot estimate quantiles below $1/N$ from a history set of N observations, with typical values for N being 260 or 520.

In this context, stress test models also have important operational requirements, akin to those for VaR models. Models that condition on recent innovations, for example the various GARCH types, yield the observed non-normality, but require daily updates to shock sizes. For reasonably complex portfolios and large collections of risk factors, such updates are impractical. They can also lead to excessive day-to-day variation in stress test results (Gençay et al, 2003, p353)².

This paper describes a stylised model based on traditional stochastic volatility approaches that is unconditional, time-homogenous, fat-tailed and easily tuneable.

2 Model Definition

Each stock, exchange rate, interest rate or other risk factor is represented by the stochastic process X_t in the stochastic volatility model

$$\begin{aligned} X_t &= \exp Y_t & dY_t &= \exp \frac{1}{2} V_t dW_t \\ dV_t &= -gV_t dt + h dW_t^{(1)} & dW_t &= \rho dW_t^{(1)} + \sqrt{1 - \rho^2} dW_t^{(2)} \end{aligned} \quad (1)$$

where $W_t^{(1)}$ and $W_t^{(2)}$ are independent Brownian motions. The model is lognormal with a correlated stochastic volatility that follows an Ornstein-Uhlenbeck process. Setting ρ and g to zero yields the stochastic volatility model of Hull and White (1987), which was generalised by Scott (1987) to incorporate mean-reverting volatility, and independently by Wiggins (1987) to incorporate correlations between price and volatility. The model incorporates both of these generalisations.

The solution for the Ornstein-Uhlenbeck volatility process is

$$\begin{aligned} V_t &= \exp(-gt) \left\{ V_0 + \int_0^t \exp(gs) dW_s^{(1)} \right\} \quad \Rightarrow \quad V_t \sim \mathbf{N}(\mu(t), \theta(t)) \\ \text{where } \mu(t) &= V_0 \exp(-gt) \rightarrow 0, \quad \theta(t) = \frac{h^2}{2g} [1 - \exp(-2gt)] \rightarrow \frac{h^2}{2g}. \end{aligned} \quad (2)$$

In the limit as time passes, the distribution for V_t becomes stationary at a normal with mean zero and variance $\frac{h^2}{2g}$. We are interested in extreme quantiles of the change in Y_t over different holding periods under this stationary distribution, that is, unconditional on the initial value of V_0 . We will calibrate the model by choosing g , h and ρ , as described later.

Note that, in applying this model to a given risk factor, Y_t may represent the absolute value, or the log returns, or even some other function of the

²Another approach to estimating extreme shocks is Extreme Value Theory (Embrechts et al., 1997); however it is not time homogenous, and requires careful judgement based on detailed data analysis of each individual risk factor to calibrate correctly.

price. The model yields quantiles of Y_t that can be converted into quantiles of any desired function, and hence shock sizes of various types (absolute, relative and so on). For the purposes of stress testing, we operate the model in the real-world measure, since it is real-world probabilities that are of interest, not risk-neutral ones.

There is no need to include a drift term other than the mean-reversion drift, $-gV_t$, in the volatility equation (as would be usual in a general Ornstein-Uhlenbeck formulation). Such a term would merely shift the stationary distribution for V_t to a different mean, leading to a fixed rescaling of Y_t . Since later we will express all shock sizes as multiples of the standard deviation, this has no impact³.

Likewise, the model contains no drift term in the spot equation. Such a term is unnecessary because shock events occur over relatively short periods. The expected daily return of a risk factor, even if it is non-zero, will be insignificant in comparison to size of the shock in a stress event. There is also no mean reversion term in the spot equation. This leads to conservative shock sizes, since any mean reversion in the spot process reduces the variance of spot over a finite period and leads to narrower quantiles.

Where scenarios shock several risk factors simultaneously, we apply the model to each separately to derive the shocks. This procedure is justified by the belief that, under stressed market conditions, normal correlation patterns break down, and is equivalent to assuming that correlations between shocked risk factors become either +1 or -1 (as required by APRA, 2000, paragraph 30). Where there is an expectation that two risk factors may not move together under stress, separate scenarios can be defined to move each one. When risk factors are moved together, they are assumed perfectly correlated, eliminating any possible diversification benefit and typically giving a conservative picture of risk.

The model has fat tails due to the stochastic volatility, which are most pronounced in the instantaneous distribution. Using the fact that V_t is normal, the instantaneous kurtosis, κ , of Y_t , follows directly from (1).

$$\kappa = \frac{\mathbf{E}[dY_t^4]}{\mathbf{E}[dY_t^2]^2} = 3 \exp \frac{h^2}{2g}$$

Since values of h lead to unique values of κ for a given g , we will usually use g , κ and ρ as the model parameters, rather than the original g , h and ρ . κ enables us to tune the degree of excess kurtosis, ρ captures the impact of asymmetry or skew, and g controls the rate of mean-reversion in volatility which leads to decay in the excess kurtosis over time.

Conservatively high shock sizes result when κ is over-estimated. Higher κ (or equivalently higher h) fattens the tails, leading to wider quantiles. Like-

³An interesting possibility is to introduce a volatility drift of $-\frac{h^2}{4}$, which results in the mean level of V_t being $-\frac{h^2}{4g}$ and the variance of Y_t over 1 time unit being unitary.

wise, conservatively high shock sizes in the upward direction result when ρ is over-estimated, since higher ρ increases the skewness of the price distribution, leading to higher quantiles in the upper tail. In the downward direction, the reverse is true. However, for stress testing, we finesse this issue by always computing the quantile on the upper side and then applying the result in both upward and downward shifts, so that shock sizes are conservative when ρ is over-estimated⁴. Shock sizes in the model respond to g non-monotonically. Generally lower g leads to wider quantiles. However, there is a region of the parameter space where the quantiles narrow mildly for decreasing g as shown below. Nonetheless, it is generally conservative to under-estimate g .

3 Computing Quantiles

We seek extreme quantiles of the change in Y_t , unconditional on V_0 . The exact distribution for Y_t is not analytically tractable, so to compute quantiles we use a moment-matching strategy. A one-period version of the model is constructed so as to have the same first four moments as the continuous model, and then quantiles are computed using the one-period version. The one-period model mimics the original continuous stochastic volatility model, as follows.

$$Y = B + e^{A+Hz_2} z_1$$

where z_1 and z_2 are standard normal with $\langle z_1 z_2 \rangle = \rho$.

Given the similar forms of the one-period model and the continuous model, and four matched moments, we can expect the quantiles of the two models to be closely aligned. However, because we are dealing with extreme quantiles, there is a possibility of divergence. To check that the difference is not excessive, we have conducted a simulation, as reported below.

So we require the first four moments of Y_t , unconditional on V_0 . Because Y_t is a martingale, whatever the correlation ρ , and using (1) we immediately have the first two moments as

$$\begin{aligned} \mathbf{E}Y_t &= 0 \quad \text{and} \\ \mathbf{E}Y_t^2 &= \mathbf{E}\langle Y \rangle_t = \int_0^t \mathbf{E} \exp V_s ds = \int_0^t \exp \left[\mu(s) + \frac{1}{2} \theta(s) \right] ds \end{aligned} \quad (3)$$

which, in the unconditional case, becomes

$$\mathbf{E}Y_t^2 = \int_0^t \exp \left(\frac{h^2}{4g} \right) ds = t \exp \left(\frac{h^2}{4g} \right). \quad (4)$$

⁴In equity and fixed interest markets, the observed correlation between log returns and volatility is typically negative. However, in others, such as commodities markets, it is commonly positive. The model is symmetric in ρ ; changing the sign of ρ simply switches the upper tail quantiles with the lower tail ones. We will therefore always work with positive values of ρ .

Appendix A derives the third and fourth moments in terms of integrals of deterministic functions. The third moment scales linearly with ρ (and is zero when ρ is zero) and involves integration over two dimensions. The fourth moment involves terms requiring integration over one, two and three dimensions.

In order to match the one-period model to the continuous model, we also require the first four moments of the one-period model. These are analytically derived in Appendix B, which also shows how to solve for the one-period model parameters so as to match moments with the continuous model. This requires numerically solving a two-dimensional polynomial equation.

Once the one-period model has been moment-matched to the continuous model for a given holding period, corresponding quantiles can be computed in the one-period model. While not analytic, these quantiles can be found by a one-dimensional numerical integration involving quantiles of the normal distribution. The expression for this integral is also given in Appendix B.

Table 1 shows the resulting quantiles expressed as numbers of standard deviations for each of seven holding periods, using a range of values for g and κ . In all cases ρ equals 0.5, which is arbitrary but turns out to have little impact.

The most important model parameter is κ , for which higher values dramatically raise the quantiles. Changes to g make a smaller difference, typically increasing quantiles over longer holding periods as g decreases (that is, as inverse mean reversion⁵ increases). Changes to ρ , although not presented here, make an even smaller difference. The non-monotonicity in g is visible in the one month holding period, although it is mild⁶. While there is no guaranteed conservative method for choosing g , erring on the low side will tend to make quantiles larger over longer holding periods and only marginally smaller over shorter ones.

To derive shock sizes for stress testing, values in Table 1 can be scaled to multiples of daily standard deviation using (4) and then multiplied by the historical daily standard deviation of changes in each risk factor.

4 Simulation Check

In order to check how accurate the moment-matching approximation is, and also to independently confirm values for the moments of the continuous model, a simulator was implemented. The design of the simulator is unusual in that spot paths are not explicitly simulated. Instead, only volatility paths

⁵The inverse mean reversion rate is sometimes referred to as the characteristic time.

⁶Note that this non-monotonicity is not due only to the moment-matching approximation. The same effect can be observed in the simulator, which estimates the quantiles of the continuous model directly.

Default Probability	Holding Period	Inverse Mean Reversion (g^{-1})					
		1m	2m	3m	4m	5m	6m
Kurtosis (κ) 7		Quantiles as Multiples of Standard Deviation					
0.00000115	1d	13.648	13.397	13.278	13.204	13.153	13.115
0.00000575	1w	11.728	11.470	11.317	11.215	11.141	11.084
0.00001150	2w	10.837	10.684	10.549	10.448	10.371	10.309
0.00002500	1m	9.622	9.730	9.672	9.601	9.536	9.480
0.00007500	3m	7.446	8.033	8.207	8.261	8.271	8.261
0.00015000	6m	6.058	6.757	7.073	7.233	7.320	7.368
0.00030000	1y	4.934	5.519	5.868	6.092	6.241	6.343
Kurtosis (κ) 10		Quantiles as Multiples of Standard Deviation					
0.00000115	1d	17.485	17.148	16.986	16.886	16.817	16.765
0.00000575	1w	14.599	14.279	14.083	13.951	13.855	13.781
0.00001150	2w	13.285	13.125	12.960	12.834	12.737	12.659
0.00002500	1m	11.551	11.747	11.693	11.613	11.537	11.469
0.00007500	3m	8.624	9.402	9.648	9.733	9.758	9.755
0.00015000	6m	6.838	7.733	8.141	8.354	8.472	8.540
0.00030000	1y	5.414	6.163	6.603	6.884	7.072	7.204
Kurtosis (κ) 13		Quantiles as Multiples of Standard Deviation					
0.00000115	1d	20.445	20.041	19.846	19.726	19.642	19.579
0.00000575	1w	16.752	16.388	16.160	16.005	15.892	15.804
0.00001150	2w	15.095	14.932	14.747	14.602	14.490	14.399
0.00002500	1m	12.957	13.215	13.167	13.081	12.997	12.921
0.00007500	3m	9.478	10.381	10.674	10.782	10.816	10.818
0.00015000	6m	7.411	8.430	8.897	9.143	9.282	9.364
0.00030000	1y	5.774	6.629	7.126	7.442	7.654	7.803
Kurtosis (κ) 16		Quantiles as Multiples of Standard Deviation					
0.00000115	1d	22.873	22.412	22.190	22.053	21.958	21.886
0.00000575	1w	18.483	18.084	17.830	17.657	17.530	17.432
0.00001150	2w	16.537	16.373	16.171	16.012	15.887	15.787
0.00002500	1m	14.066	14.373	14.329	14.239	14.148	14.065
0.00007500	3m	10.150	11.144	11.472	11.595	11.638	11.643
0.00015000	6m	7.867	8.974	9.482	9.751	9.905	9.996
0.00030000	1y	6.065	6.998	7.532	7.872	8.101	8.261

Table 1: Quantile Results for Given Default Probabilities and Holding Periods (d=days; w=weeks; m=months) at Various Levels of κ and g

are generated, and numerical methods are used to compute the quantiles of the spot distribution via integrals over the volatility paths.

This is possible because the distribution of Y_t , conditioned on the filtration generated by $W_t^{(1)}$, is normal. The unconditional distribution of Y_t is therefore a continuous mixture of normals, with the weightings in the mixture being determined by the joint distribution of the mean and variance of Y_t conditioned on $W_s^{(1)}$ for $0 \leq s \leq t$. The simulator builds a sample of this joint distribution containing one point for each volatility path that it generates. Cumulative densities for Y_t can then be computed as a weighted sum of the mixture densities, and quantiles can be found by numerical inversion of the cumulative densities. Full details are given in Appendix C. Appendix C also derives formulae for the first four moments of Y_t as weighted combinations of the mixture moments, thus allowing the simulator to estimate the moments of Y_t , which yields a means of checking the moments obtained using the formulae derived in Appendix A.

Simulation runs have been performed for various cases, including $g = 4$ (inverse mean reversion of 3 months), $\kappa = 13$ and $\rho = 0.5$ for two holding periods. Over a one year holding period, the result in this case was 7.1311 ± 0.0054 which includes the value 7.126 found using moment-matching and numerical integration. Over a one day holding period, the result was 18.6763 ± 0.1748 which is about 6% narrower than the 19.846 found using moment-matching and numerical integration. The approximation arising from moment-matching appears to be small, and probably conservative.

Also, the moments calculated by the simulator match closely those found by numerical integration.

5 Calibration

It remains to choose settings for g , κ and ρ . We briefly present two methods below. However, rather than attempting to calibrate the model for every risk factor individually, our approach is to fit calibrations to a broad range of relevant risk factors and then choose a fixed parameter setting that yields conservative shocks for all those risk factors (or at least, all risk factors of a given type). Thereafter, historical data for individual risk factors is used only to determine the daily standard deviation, of which the shocks sizes are fixed multiples as determined by the chosen parameter setting. We do this because statistically precise estimates of the model parameters require large amounts of historical data and only risk factors with long history sets could be accurately calibrated on an individual basis. Our calibration methods are thus only intended to indicate the range of reasonable parameter values, from which we can select conservative settings.

Our first method relies on extracting information from the historical




Figure 1: Volatility of Implied Vols ($g = 4yr^{-1}$; $\kappa = 4,7,10,13,16$)




Figure 2: Volatility of Implied Vols ($g = 1,2,4,6,12yr^{-1}$; $\kappa = 13$)

volatility of implied volatilities, and thus requires a liquid market in options on the underlying risk factor. It is based on how the model implies that $\mathbf{E}Y_t^2$ will vary (under the real-world measure). If we assume that the square of the at-the-money implied volatility in options prices, σ_{BS}^2 , is equal to $\mathbf{E}Y_t^2$, then the model predicts how these implied volatilities should vary. Although

implied volatilities are characteristics of the risk-neutral rather than the real-world measure, the two are equal under the Black-Scholes assumptions, and should be close for realistic violations of those assumptions⁷.

Details are presented in Appendix D showing that g and h determine the full term structure of volatility of (the log of) implied volatilities. Figures 1 and 2 show this term structure for various values of g and κ .

As the graphs show κ drives the volatility of implied vols across the whole term structure, while g determines the rate at which those volatilities fall as expiry increases. Thus, given the historical volatility of implied volatility for two maturities, we can fit values for g and κ to match. If we have three or more maturities, we can fit the entire term structure by, for example, least squares.

To calibrate ρ , we consider what the model implies about the correlation between $\mathbf{E}Y_t^2$ and the real-world changes in Y_t . Making the same assumption about σ_{BS}^2 , this yields a correlation between changes in Y_t and changes in (the log of) at-the-money implied volatilities. Appendix D gives this correlation as a function of g , h , and ρ that is linear in ρ . It is therefore simple, given values for g and h , to determine a value for ρ to match historical correlations of Y_t and at-the-money implied volatilities for Y_t .




Figure 3: Probability Distribution of Model Compared to Normal (dotted) Over One Day ($g = 4yr^{-1}$; $\kappa = 10$, $\rho = 0.5$)

Our second method is based on inter-quartile ranges. Figure 3 shows that the distribution of returns given by the model has fat tails relative to the normal. However, to compensate for this, it has thinner shoulders. This is especially pronounced for inter-quartile ranges of 25-35% on either side of

⁷Note that the use of ATM implied volatilities ignores market skew. However, incorporating skew information would require modelling the real-world drift and converting between it and the risk-neutral drift.

the mean.

The shoulders get thinner as κ increases, so it is possible to calibrate the model using inter-quartile ranges (IQR). Unlike the extreme quantiles, inter-quartile ranges can be estimated with reasonable statistical confidence from much smaller history sets. Figure 4 shows how the statistic IQR/SD varies for the model distribution as κ varies.




Figure 4: Daily Inter-quartile Range As κ Varies ($g = 4yr^{-1}$; $\rho = 0.5$)

For kurtosis values below about 20, the steep slope of this curve means that a given IQR/SD value will accurately discriminate the corresponding kurtosis. The IQR/SD statistic turns out to be extremely insensitive to the other model parameters, g and ρ , so if we assume that data are distributed according to the model, we can accurately determine κ from IQR/SD by inverting the function in Figure 4. Since κ is the most important parameter, this yields a method of calibration, so long as we are prepared to choose values for g and ρ .

Of course, any estimate of the IQR/SD statistic is subject to sampling uncertainty, which could make a significant difference to the resulting κ . Although the same argument applies to the standard deviation itself, and hence to the shock sizes, the confidence bounds on standard deviation estimates are narrower than those on the IQR/SD statistic. A conservative approach is therefore to construct a confidence interval for the IQR/SD statistic, then take the resulting lower bound instead of the central estimate⁸.

⁸Large sample confidence bounds for IQR/SD are straight-forward to construct using the asymptotic distributions of its components and conservative estimates for their correlations.

References

- [1] Alexander, C. and Sheedy, E. 2007 Model-Based Stress Tests: Linking Stress Tests to VaR for Market Risk. International Capital Markets Association Centre Discussion Papers in Finance DP2007-02
- [2] Australian Prudential Regulation Authority (APRA) 2000. Guidance Note AGN 113.2 – The Internal Model Approach.
- [3] Berkowitz, J. 1999. A Coherent Framework for Stress-testing. Finance and Economics Discussion Series 1999-29, Board of Governors of the Federal Reserve System (U.S.).
- [4] Committee on the Global Financial System 2005. A Survey of Stress Tests and Current Practice at Major Financial Institutions.
- [5] Embrechts, P., Klüppelberg, C. and Mikosch, T. 1997. Modelling Extremal Events for Insurance and Finance. Springer, Berlin.
- [6] Gençay, R., Selçuk, F. and Ulugülyağci, A. 2003. High Volatility, Thick Tails and Extreme Value Theory in Value-at-risk Estimation. In *Insurance: Mathematics and Economics*, Vol. 33, issue 2, pp337-356.
- [7] Hull, J. and White, A. 1987. The Pricing of Options on Assets with Stochastic Volatilities. In *Journal Of Finance*, Vol. 42, pp281-300.
- [8] Scott, L. 1987. Option Pricing When the Variance Changes Randomly – Theory, Estimation, and an Application. In *Journal Of Financial And Quantitative Analysis*, Vol. 22, No. 4, pp419-438.
- [9] Wiggins, J. 1987. Option Values Under Stochastic Volatility – Theory and Empirical Estimates. In *Journal Of Financial Economics*, Vol. 19, pp351-372.

A Continuous Model Moments

A.1 Introduction

Required are the first four moments of Y_t . We begin by considering the distribution conditional on V_0 .

Since V_t is a mean-reverting Ornstein Uhlenbeck process with solution given by (2), the covariance $C(s, u)$ of V_s and V_u is

$$C(s, u) = \text{cov } V_s V_u = \frac{h^2}{2g} \{ \exp(-g|u-s|) - \exp(-g[u+s]) \}$$

$$\text{with } C(s, s) = \theta(s) \quad \text{and} \quad \frac{\partial C(s, u)}{\partial u} = -gC(s, u) \quad \text{for } s \leq u.$$

From this, together with (2), it is easy to derive the following useful expected value formulæ:

1. For any constants a, b, c and times s, u, v

$$\begin{aligned} \mathbf{E} \exp[aV_s + bV_u + cV_v] &= F(a, b, c; s, u, v) \\ &= \exp \left[\begin{array}{c} a\mu(s) + b\mu(u) + c\mu(v) \\ + \frac{1}{2}a^2\theta(s) + \frac{1}{2}b^2\theta(u) + \frac{1}{2}c^2\theta(v) \\ ab C(s, u) + ac C(s, v) + bc C(u, v) \end{array} \right], \end{aligned}$$

2. Differentiating partially with respect to a

$$\begin{aligned} \mathbf{E} V_s \exp[aV_s + bV_u + cV_v] &= F_a(a, b, c, s, u, v) \\ &= [\mu(s) + a\theta(s) + b C(s, u) + c C(s, v)] F(a, b, c, s, u, v), \end{aligned}$$

3. And differentiating partially with respect to a and b

$$\begin{aligned} F_{a,b}(a, b, c, s, u, v) &= \mathbf{E} V_s V_u \exp[aV_s + bV_u + cV_v] = \\ &\left\{ \begin{array}{c} [\mu(s) + a\theta(s) + b C(s, u) + c C(s, v)] \\ \times [\mu(u) + b\theta(u) + a C(s, u) + c C(u, v)] \\ + C(s, u) \end{array} \right\} F(a, b, c, s, u, v). \end{aligned}$$

If we now introduce martingale variables $Y_t^{(1)}$ and $Y_t^{(2)}$ defined by

$$Y_t^{(1)} = \int_0^t \exp \frac{1}{2} V_s dW_s^{(1)}, \quad Y_t^{(2)} = \int_0^t \exp \frac{1}{2} V_s dW_s^{(2)} \quad \Rightarrow \quad Y_t = \rho Y_t^{(1)} + \sqrt{1 - \rho^2} Y_t^{(2)}$$

and $\mathcal{F}_t^{(1)}$ is the filtration generated by $W_t^{(1)}$, then $Y_t^{(1)}$ is $\mathcal{F}_t^{(1)}$ -measurable. So applying Ito to $\exp \frac{1}{2} V_t$ connects $Y_t^{(1)}$ to V_t via some *finite variation* integrals

$$\frac{h}{2} Y_t^{(1)} = \exp \frac{1}{2} V_t - \exp \frac{1}{2} V_0 + \frac{g}{2} \int_0^t V_s \exp \frac{1}{2} V_s ds - \frac{h^2}{8} \int_0^t \exp \frac{1}{2} V_s ds. \quad (5)$$

Next, applying Ito to Y_t^3 and Y_t^4 , using the martingale property of Y_t , and integrating by parts, produces the following useful formulæ for the third and fourth moments

$$\begin{aligned} \mathbf{E}Y_t^3 &= 3\mathbf{E} \int_0^t Y_s \exp V_s ds = 3\mathbf{E}Y_t \int_0^t \exp V_s ds, \\ \mathbf{E}Y_t^4 &= 6\mathbf{E} \int_0^t Y_s^2 \exp V_s ds = 6\mathbf{E} \left\{ Y_t^2 \int_0^t \exp V_s ds - \int_0^t \left(\int_0^s \exp V_u du \right) \exp V_s ds \right\}, \\ &= 6\mathbf{E}Y_t^2 \int_0^t \exp V_s ds - 3\mathbf{E} \left(\int_0^t \exp V_s ds \right)^2. \end{aligned} \quad (6)$$

In the *uncorrelated* case $\rho = 0$ conditioning on $\mathcal{F}_t^{(1)}$ in these formulæ gives

$$\begin{aligned} \mathbf{E}Y_t^3 &= 3\mathbf{E} \left\{ \int_0^t \exp V_s ds \mathbf{E} \left\{ Y_t^{(2)} \middle| \mathcal{F}_t^{(1)} \right\} \right\} = 0 \\ \mathbf{E}Y_t^4 &= 3\mathbf{E} \left(\int_0^t \exp V_s ds \right)^2 = 3 \int_0^t \int_0^t \mathbf{E} \exp (V_s + V_u) ds du, \\ &= 3 \int_0^t \int_0^t F(1, 1, 0, s, u, 0) ds du. \end{aligned} \quad (7)$$

In both the uncorrelated and correlated cases the volatility process V_t is stationary, so that after running for some time the initial conditions *wash out* and V_t becomes normally distributed with constant variance $V_t \sim \mathbf{N} \left(0, \frac{h^2}{2g} \right)$. So assume both Y_t and V_t have been running for a long time in the past and at time $t = 0$ have settled into *pseudo-stationary mode* (i.e V_t stationary, but not of course Y_t). In this unconditional case V_0 has an initial distribution $V_0 \sim \mathbf{N} \left(0, \frac{h^2}{2g} \right)$ and

$$\begin{aligned} G(a, b, c; s, u, v) &= \mathbf{E} \exp [aV_s + bV_u + cV_v] \\ &= \mathbf{E}^{V_0} \{ \mathbf{E} \exp [aV_s + bV_u + cV_v] | V_0 \} = \mathbf{E}^{V_0} F(a, b, c; s, u, v). \\ &= \exp \left[\frac{h^2}{4g} \left(a^2 + b^2 + c^2 + 2abe^{-g|u-s|} + 2ace^{-g|v-s|} + 2bce^{-g|u-v|} \right) \right] \end{aligned}$$

Hence in the unconditional case, replacing $F(\cdot)$ by $G(\cdot)$:

1. For any constants a, b, c and times s, u, v

$$\mathbf{E} \exp [aV_s + bV_u + cV_v] = G(a, b, c; s, u, v),$$

2. Differentiating partially with respect to a

$$\begin{aligned} \mathbf{E}V_s \exp [aV_s + bV_u + cV_v] &= G_a(a, b, c; s, u, v) \\ &= \frac{h^2}{2g} \left(a + be^{-g|u-s|} + ce^{-g|v-s|} \right) G(a, b, c; s, u, v), \end{aligned}$$

3. And differentiating partially with respect to a and b

$$\begin{aligned} \mathbf{E}V_s V_u \exp[aV_s + bV_u + cV_v] &= G_{a,b}(a, b, c; s, u, v) \\ &= \left\{ \frac{h^4}{4g^2} \left(a + b\mathbf{e}^{-g|u-s|} + c\mathbf{e}^{-g|v-s|} \right) \left(b + a\mathbf{e}^{-g|u-s|} + c\mathbf{e}^{-g|u-v|} \right) + \frac{h^2}{2g} \mathbf{e}^{-g|u-s|} \right\} \\ &\quad \times G(a, b, c; s, u, v). \end{aligned}$$

If some of the a, b or c are zero, with obvious notation, we truncate and write

$$\begin{aligned} G(a; s) &= \exp \frac{h^2}{4g} a^2 & G_a(a; s) &= \frac{h^2}{2g} a \exp \frac{h^2}{4g} a^2 \\ G(a, b; s, u) &= \exp \left[\frac{h^2}{4g} \left(a^2 + b^2 + 2ab\mathbf{e}^{-g|u-s|} \right) \right] \\ G_a(a, b; s, u) &= \frac{h^2}{2g} \left(a + b\mathbf{e}^{-g|u-s|} \right) G(a, b; s, u) \\ G_{a,b}(a, b; s, u) &= \left\{ \frac{h^4}{4g^2} \left(a + b\mathbf{e}^{-g|u-s|} \right) \left(b + a\mathbf{e}^{-g|u-s|} \right) + \frac{h^2}{2g} \mathbf{e}^{-g|u-s|} \right\} G(a, b; s, u). \end{aligned}$$

So for the unconditional uncorrelated case, replacing $F(\cdot)$ by $G(\cdot)$ in (7) we have

$$\mathbf{E}Y_t^4 = 3 \int_0^t \int_0^t G(1, 1, 0; s, u, 0) dsdu = 3 \exp \frac{h^2}{2g} \int_0^t \int_0^t \exp \left[\frac{h^2}{2g} \exp(-g|u-s|) \right] dsdu.$$

A.2 Third Moment in the Correlated Case

Condition on $\mathcal{F}_t^{(1)}$ in (6) and substitute $Y_t^{(1)}$ from (5) to get

$$\begin{aligned} \mathbf{E}Y_t^3 &= 3\mathbf{E} \left[\rho Y_t^{(1)} + \sqrt{1-\rho^2} Y_t^{(2)} \right] \int_0^t \exp V_u du = 3\rho \mathbf{E}Y_t^{(1)} \int_0^t \exp V_u du, \\ &= \frac{6\rho}{h} \mathbf{E} \left\{ \left[\frac{\exp \frac{1}{2} V_t - \exp \frac{1}{2} V_0}{+ \frac{g}{2} \int_0^t V_s \exp \frac{1}{2} V_s ds - \frac{h^2}{8} \int_0^t \exp \frac{1}{2} V_s ds} \right] \times \int_0^t \exp V_u du \right\}, \\ &= \frac{6\rho}{h} \int_0^t F\left(1, \frac{1}{2}; u, t\right) du - \frac{6\rho}{h} \int_0^t F\left(1, \frac{1}{2}; u, 0\right) du \\ &\quad + \frac{3\rho g}{h} \int_0^t \int_0^t F_a\left(\frac{1}{2}, 1; s, u\right) dsdu - \frac{3\rho h}{4} \int_0^t \int_0^t F\left(\frac{1}{2}, 1; s, u\right) dsdu. \end{aligned}$$

In the unconditional case, replacing $F(\cdot)$ by $G(\cdot)$ the first two terms cancel because

$$\begin{aligned} \int_0^t G\left(1, \frac{1}{2}; u, t\right) du &= \int_0^t \exp \left[\frac{h^2}{16g} \left(5 + 4\mathbf{e}^{-g[t-u]} \right) \right] du \\ &= \int_0^t \exp \left[\frac{h^2}{16g} \left(5 + 4\mathbf{e}^{-gu} \right) \right] du = \int_0^t G\left(1, \frac{1}{2}; u, 0\right) du, \end{aligned}$$

and asymptotically the second two become

$$\begin{aligned} 3\rho \int_0^t \int_0^t \left\{ \frac{g}{h} G_a \left(\frac{1}{2}, 1; s, u \right) - \frac{h}{4} G \left(\frac{1}{2}, 1; s, u \right) \right\} ds du, \\ = 3\rho h \int_0^t \int_0^u e^{-g[u-s]} G \left(\frac{1}{2}, 1; s, u \right) ds du. \end{aligned}$$

A.3 Fourth Moment in the Correlated Case

In (6) expand Y_t^2 so as to be able to condition on $\mathcal{F}_t^{(1)}$ in

$$\begin{aligned} \mathbf{E} Y_t^2 \int_0^t \exp V_s ds &= \mathbf{E} \left[\begin{array}{c} \rho^2 \left(Y_t^{(1)} \right)^2 + (1 - \rho^2) \left(Y_t^{(2)} \right)^2 \\ + 2\rho \sqrt{1 - \rho^2} Y_t^{(1)} Y_t^{(2)} \end{array} \right] \int_0^t \exp V_u du, \\ &= \rho^2 \mathbf{E} \left(Y_t^{(1)} \right)^2 \int_0^t \exp V_u du + (1 - \rho^2) \mathbf{E} \left(\int_0^t \exp V_u du \right)^2. \end{aligned}$$

Hence, after substituting back in (6)

$$\mathbf{E} Y_t^4 = 6\rho^2 \mathbf{E} \left(Y_t^{(1)} \right)^2 \int_0^t \exp V_u du + 3(1 - 2\rho^2) \mathbf{E} \left(\int_0^t \exp V_s ds \right)^2, \quad (8)$$

an expression which tallies with the uncorrelated result (7) when $\rho = 0$.

Substituting from (5), the 4th moment is therefore

$$\begin{aligned} \mathbf{E} Y_t^4 &= 24 \frac{\rho^2}{h^2} \mathbf{E} \left\{ \left[\begin{array}{c} \exp \frac{1}{2} V_t - \exp \frac{1}{2} V_0 \\ + \frac{g}{2} \int_0^t V_s \exp \frac{1}{2} V_s ds - \frac{h^2}{8} \int_0^t \exp \frac{1}{2} V_s ds \end{array} \right]^2 \int_0^t \exp V_v dv \right\} \\ &\quad + 3(1 - 2\rho^2) \int_0^t \int_0^t F(1, 1, 0; s, u, 0) ds du. \end{aligned} \quad (9)$$

Grouping according to $F(\cdot)$ functions, the terms in the expansion of the $\mathbf{E}\{\cdot\}$ are

$$\begin{aligned} \mathbf{E} \left[\exp \frac{1}{2} V_0 \right]^2 \int_0^t \exp V_v dv &= \int_0^t F \left(\frac{1}{2}, \frac{1}{2}, 1; 0, 0, v \right) dv \\ \mathbf{E} \left[\exp \frac{1}{2} V_t \right]^2 \int_0^t \exp V_v dv &= \int_0^t F \left(\frac{1}{2}, \frac{1}{2}, 1; t, t, v \right) dv \\ \mathbf{E} \left[-2 \exp \frac{1}{2} V_t \exp \frac{1}{2} V_0 \right] \int_0^t \exp V_v dv &= -2 \int_0^t F \left(\frac{1}{2}, \frac{1}{2}, 1; t, 0, v \right) dv \\ \mathbf{E} \left[g \exp \frac{1}{2} V_t \int_0^t V_s \exp \frac{1}{2} V_s ds \right] \int_0^t \exp V_v dv &= g \int_0^t \int_0^t F_a \left(\frac{1}{2}, \frac{1}{2}, 1; s, t, v \right) ds dv \\ \mathbf{E} \left[-\frac{h^2}{4} \exp \frac{1}{2} V_t \int_0^t \exp \frac{1}{2} V_s ds \right] \int_0^t \exp V_v dv &= -\frac{h^2}{4} \int_0^t \int_0^t F \left(\frac{1}{2}, \frac{1}{2}, 1; s, t, v \right) ds dv \end{aligned}$$

$$\begin{aligned}\mathbf{E} \left[-g \exp \frac{1}{2} V_0 \int_0^t V_s \exp \frac{1}{2} V_s ds \right] \int_0^t \exp V_v dv &= -g \int_0^t \int_0^t F_a \left(\frac{1}{2}, \frac{1}{2}, 1; s, 0, v \right) ds dv \\ \mathbf{E} \left[\frac{h^2}{4} \exp \frac{1}{2} V_0 \int_0^t \exp \frac{1}{2} V_s ds \right] \int_0^t \exp V_v dv &= \frac{h^2}{4} \int_0^t \int_0^t F \left(\frac{1}{2}, \frac{1}{2}, 1; s, 0, v \right) ds dv\end{aligned}$$

$$\begin{aligned}\mathbf{E} \left[\frac{g}{2} \int_0^t V_s \exp \frac{1}{2} V_s ds \right]^2 \int_0^t \exp V_v dv &= \frac{g^2}{4} \int_0^t \int_0^t \int_0^t F_{a,b} \left(\frac{1}{2}, \frac{1}{2}, 1; s, u, v \right) ds dv du \\ \mathbf{E} \left[\frac{h^2}{8} \int_0^t \exp \frac{1}{2} V_s ds \right]^2 \int_0^t \exp V_v dv &= \frac{h^4}{64} \int_0^t \int_0^t \int_0^t F \left(\frac{1}{2}, \frac{1}{2}, 1; s, u, v \right) ds dv du \\ \mathbf{E} \left[\begin{aligned} & -\frac{gh^2}{8} \int_0^t V_s \exp \frac{1}{2} V_s ds \\ & \times \int_0^t \exp \frac{1}{2} V_s ds \end{aligned} \right] \int_0^t \exp V_v dv &= -\frac{gh^2}{8} \int_0^t \int_0^t \int_0^t F_a \left(\frac{1}{2}, \frac{1}{2}, 1; s, u, v \right) ds dv du\end{aligned}$$

For the unconditional case replace $F(\cdot)$ by $G(\cdot)$. The single integral terms become

$$\begin{aligned}& \int_0^t \left\{ G \left(\frac{1}{2}, \frac{1}{2}, 1; 0, 0, v \right) - 2G \left(\frac{1}{2}, \frac{1}{2}, 1; t, 0, v \right) + G \left(\frac{1}{2}, \frac{1}{2}, 1; t, t, v \right) \right\} dv \\ &= \int_0^t \left\{ 2 \exp \left[\frac{h^2}{2g} (1 + \mathbf{e}^{-gv}) \right] - 2 \exp \left[\frac{h^2}{8g} (3 + \mathbf{e}^{-gt} + 2\mathbf{e}^{-g[t-v]} + 2\mathbf{e}^{-gv}) \right] \right\} dv.\end{aligned}$$

The double integral terms together cancel, because setting

$$\begin{aligned}f(s, v) &= gG_a \left(\frac{1}{2}, \frac{1}{2}, 1; s, 0, v \right) + \frac{h^2}{4} G \left(\frac{1}{2}, \frac{1}{2}, 1; s, 0, v \right), \\ \int_0^t \int_0^t & \left\{ \begin{aligned} & gG_a \left(\frac{1}{2}, \frac{1}{2}, 1; s, t, v \right) - \frac{h^2}{4} G \left(\frac{1}{2}, \frac{1}{2}, 1; s, t, v \right) \\ & -gG_a \left(\frac{1}{2}, \frac{1}{2}, 1; s, 0, v \right) + \frac{h^2}{4} G \left(\frac{1}{2}, \frac{1}{2}, 1; s, 0, v \right) \end{aligned} \right\} ds dv \\ &= \int_0^t \int_0^t \{ f(t-s, t-v) - f(s, v) \} ds dv = 0.\end{aligned}$$

Using the symmetry of s and u , the treble integral terms are

$$\begin{aligned}& \int_0^t \int_0^t \int_0^t \left\{ \frac{g^2}{4} G_{a,b} - \frac{gh^2}{8} G_a + \frac{h^4}{64} G \right\} \left(\frac{1}{2}, \frac{1}{2}, 1; s, u, v \right) ds dv du \\ &= \frac{h^2}{8} \int_0^t \int_0^t \int_0^t G \left(\frac{1}{2}, \frac{1}{2}, 1; s, u, v \right) \left[\begin{aligned} & \frac{h^2}{8} \mathbf{e}^{-2g|s-u|} + g\mathbf{e}^{-g|s-u|} \\ & + \frac{h^2}{2} \mathbf{e}^{-g|s-u|} \mathbf{e}^{-g|u-v|} + \frac{h^2}{2} \mathbf{e}^{-g|s-v|} \mathbf{e}^{-g|u-v|} \end{aligned} \right] ds dv du.\end{aligned}\tag{10}$$

B One-Period Model Results

We seek analytical expressions for the first four moments of the one-period model

$$Y = B + e^{A+Hz_2} z_1$$

where z_1 and z_2 are standard normal with $\langle z_1 z_2 \rangle = \rho$.

Setting $z_2 = \rho z_1 + \sqrt{1-\rho^2} z_3$ and applying the binomial theorem we have

$$Y^n = \sum_{i=0}^n \binom{n}{i} B^{n-i} e^{iA} e^{iH\sqrt{1-\rho^2}z_3} e^{iH\rho z_1} z_1^i$$

Since z_1 and z_3 are independent, we can apply expectations to each separately. We use the following relation which holds for any scalar m and standard normal variable z .

$$\mathbf{E} [z^i e^{mz}] = e^{\frac{m^2}{2}} \mathbf{E} [(m+z)^i]$$

Applying this to z_1 and z_3 yields a general formula for the n^{th} moment.

$$\mathbf{E} [Y^n] = \sum_{i=0}^n \binom{n}{i} B^{n-i} e^{\frac{H^2 i^2}{2} + Ai} \sum_{j=0}^i \binom{i}{j} \mathbf{E} [z_1^{i-j}] i^j H^j \rho^j$$

The first four moments are then given by the following formulæ.

$$\mathbf{E} [Y] = B + e^{\frac{H^2}{2} + A} H \rho \quad (11)$$

$$\mathbf{E} [Y^2] = B^2 + 2e^{\frac{H^2}{2} + A} H \rho B + e^{2H^2 + 2A} (4H^2 \rho^2 + 1) \quad (12)$$

$$\begin{aligned} \mathbf{E} [Y^3] = & B^3 + 3e^{\frac{H^2}{2} + A} H \rho B^2 + 3e^{2H^2 + 2A} (4H^2 \rho^2 + 1) B \\ & + e^{\frac{9H^2}{2} + 3A} (27H^3 \rho^3 + 9H \rho) \end{aligned} \quad (13)$$

$$\begin{aligned} \mathbf{E} [Y^4] = & B^4 + 4e^{\frac{H^2}{2} + A} H \rho B^3 + 6e^{2H^2 + 2A} (4H^2 \rho^2 + 1) B^2 \\ & + 4e^{\frac{9H^2}{2} + 3A} (27H^3 \rho^3 + 9H \rho) B + e^{8H^2 + 4A} (256H^4 \rho^4 + 96H^2 \rho^2 + 3) \end{aligned} \quad (14)$$

Given the values of the first four moments of the continuous model, M_1 , M_2 , M_3 , M_4 , we wish to solve for the one-period model parameters. Since M_1 is always zero, we can eliminate B using (11).

$$B = -e^{\frac{H^2}{2} + A} H \rho$$

After substituting for B , we can eliminate A using (12).

$$A = -\frac{1}{2} \log \left(\frac{-e^{H^2} H^2 \rho^2 + 4e^{2H^2} H^2 \rho^2 + e^{2H^2}}{M_2} \right)$$




Figure 5: Contours of Equations for Third and Fourth Moment

Substituting for A and B in (13) yields

$$\mathbf{E}[Y^3] = \frac{e^{\frac{3H^2}{2}} H \rho \left(2H^2 \rho^2 + 9e^{3H^2} (3H^2 \rho^2 + 1) - 3e^{H^2} (4H^2 \rho^2 + 1) \right)}{\left(\frac{e^{H^2} (e^{H^2} (4H^2 \rho^2 + 1) - H^2 \rho^2)}{M_2} \right)^{3/2}}$$

which we can simplify by scaling relative to $M_2^{3/2}$ and substituting $k = \rho H$ and $l = e^{H^2}$.

$$\frac{\mathbf{E}[Y^3]}{M_2^{3/2}} = \frac{k (9(3k^2 + 1)l^3 - 3(4k^2 + 1)l + 2k^2)}{((4k^2 + 1)l - k^2)^{3/2}} \quad (15)$$

Similarly (14) when scaled relative to M_2^2 simplifies to

$$\frac{\mathbf{E}[Y^4]}{M_2^2} = \frac{(256k^4 + 96k^2 + 3)l^6 - 36k^2(3k^2 + 1)l^3 + 6k^2(4k^2 + 1)l - 3k^4}{(k^2 - (4k^2 + 1)l)^2} \quad (16)$$

Since we seek the simultaneous solution of (15) and (16), it is illuminating to examine the contours of each. Figure 5 shows the contours for a range of

values for k and l . The contours of (15) are in red, those of (16) in black. It appears from the graph that there is at most one solution for any given set of moments.

Because $l = e^{H^2}$ and $k = \rho H$ we have $l > e^{k^2} \geq 1$ and so the denominator in (15) is always positive. Similarly the numerator must have the same sign as k , and k only appears in (16) raised to even powers. As a result, we can square (15) without changing the set of solutions, apart from the sign of k . After solving, we can recover the sign of k from the sign of M_3 .

Squaring (15) and rearranging both equations, yields the following two simultaneous polynomial equations which can be solved numerically. Note that only solutions with positive k , and where $e^{k^2} < l$, are relevant.

$$\frac{M_3^2 ((4k^2 + 1)l - k^2)^3}{M_2^3} = k^2 (9 (3k^2 + 1)l^3 - 3(4k^2 + 1)l + 2k^2)^2$$

$$\frac{M_4 (k^2 - (4k^2 + 1)l)^2}{M_2^2} = (256k^4 + 96k^2 + 3)l^6 - 36k^2(3k^2 + 1)l^3 + 6k^2(4k^2 + 1)l - 3k^4$$

Finally, given matched model parameters, we need to compute quantiles of the one-period model. We can do this by numerically inverting the cumulative density function, which we can find by integrating the region of the bivariate normal lying beyond the contour implied by the equation $B + e^{A+Hz_2}z_1 = q$. Note that $\text{Sign}(z_1) = \text{Sign}(q - B)$ so the contour lies entirely within either $z_1 > 0$ or $z_1 < 0$ depending on whether q is larger or smaller than B . Solving for the value of z_3 that makes $Y = q$ we get

$$z_3 = \frac{-A + \log\left(\frac{q-B}{z_1}\right) - Hz_1\rho}{H\sqrt{1-\rho^2}}$$

The cumulative density is therefore given by

$$\text{cum}(q) = \begin{cases} q < B & \int_{-\infty}^0 \Phi\left(-\frac{-A + \log\left(\frac{q-B}{z_1}\right) - H\rho z_1}{H\sqrt{1-\rho^2}}\right) \phi(z_1) dz_1 \\ q = B & \frac{1}{2} \\ q > B & \int_0^{\infty} \Phi\left(\frac{-A + \log\left(\frac{q-B}{z_1}\right) - H\rho z_1}{H\sqrt{1-\rho^2}}\right) \phi(z_1) dz_1 + \frac{1}{2} \end{cases}$$

C Simulator Formulæ

C.1 Simulated Quantiles

We seek a Monte Carlo estimator for quantiles of Y_t . Recall from Appendix A that

$$Y_t^{(2)} = \int_0^t e^{\frac{V_s}{2}} dW_s^{(2)}$$

where $dW_s^{(2)}$ is Brownian motion so that, conditioning on $\mathcal{F}_t^{(1)}$, $Y_t^{(2)}$ is normal with zero mean. The conditional mean and variance of Y_t are then

$$\mu = \mathbf{E} \left[Y_t | \mathcal{F}_t^{(1)} \right] = \rho Y_t^{(1)} = \rho \int_0^t e^{\frac{V_s}{2}} dW_s^{(1)} \quad \text{and}$$

$$\sigma^2 = \mathbf{E} \left[(Y_t - \mu)^2 | \mathcal{F}_t^{(1)} \right] = \mathbf{E} \left[\left(\sqrt{1 - \rho^2} Y_t^{(2)} \right)^2 \right] = (1 - \rho^2) \int_0^t e^{V_s} ds.$$

Let $p(\mu, \sigma)$ be the joint probability density of these values under $\mathcal{F}_t^{(1)}$. The probability density of Y_t is then

$$p(y) = \int \int_{-\infty}^{\infty} \phi \left(\frac{y - u}{s} \right) p(u, s) du ds. \quad (17)$$

Since the cumulative density, $c(y)$, of Y_t is simply the integral of $p(y)$ on the interval $(-\infty, y)$, and the integrations commute, we have

$$c(y) = \int \int_{-\infty}^{\infty} \Phi \left(\frac{y - u}{s} \right) p(u, s) du ds.$$

A maximum likelihood Monte Carlo estimate of $c(y)$ is therefore

$$\hat{c}(y) = \frac{1}{N} \sum_{i=1}^N \Phi \left(\frac{y - \hat{\mu}_i}{\hat{\sigma}_i} \right)$$

where the $\hat{\mu}_i$ and $\hat{\sigma}_i$ are computed from sample paths of V_t . Importantly, only two real-valued variables need to be retained for each path (and these can be computed as the path is generated so that memory requirements are constant regardless of the path length).

Since $\hat{c}(y)$ is monotonic, any arbitrary quantile, q , can be found by step search and binary chop to solve $q = \hat{c}(y)$. Note that this estimator is biased, as are all maximum likelihood quantile estimators. Therefore, care must be taken to choose sample sizes sufficiently large so as to make the bias negligible.

C.2 Simulated Moments

It is also useful for the simulator to compute moments as a check of the results in Appendix A. Like the cumulative density, non-central moments of Y_t are integrals over $p(y)$ on the real line, and again these commute with the integrals in (17). It follows that

$$\mathbf{E}[Y_t^n] = \iint_{-\infty}^{\infty} \mathbf{E}[Y_t^n | \mu=u, \sigma=s] p(u, s) du ds.$$

A Monte Carlo estimate of the n^{th} non-central moment is therefore

$$\begin{aligned} \hat{M}_n &= \frac{1}{N} \sum_{i=1}^N \mathbf{E}[Y_t^n | \mu=\hat{\mu}_i, \sigma=\hat{\sigma}_i] \\ &= \frac{1}{N} \sum_{i=1}^N \mathbf{E}[(\hat{\mu}_i + \hat{\sigma}_i z)^n] \end{aligned}$$

where z is standard normal. Estimators for the first four moments are therefore as follows.

$$\begin{aligned} \hat{M}_1 &= \frac{1}{N} \sum_{i=1}^N \hat{\mu}_i \\ \hat{M}_2 &= \frac{1}{N} \sum_{i=1}^N \hat{\mu}_i^2 + \hat{\sigma}_i^2 \\ \hat{M}_3 &= \frac{1}{N} \sum_{i=1}^N \hat{\mu}_i^3 + 3\hat{\mu}_i \hat{\sigma}_i^2 \\ \hat{M}_4 &= \frac{1}{N} \sum_{i=1}^N \hat{\mu}_i^4 + 6\hat{\mu}_i^2 \hat{\sigma}_i^2 + 3\hat{\sigma}_i^4 \end{aligned}$$

D Calibration Formulæ

We seek the unconditional instantaneous variance of log implied volatility under the continuous model, that is the variance of changes in

$$G_T(V_t) = \frac{1}{2} \log (\mathbf{E} [Y_T^2])$$

Applying Ito to $G_T(V_t)$ in the model yields

$$dG_T(V_t) = \left(\frac{\partial G}{\partial V}(-gV_t) + \frac{1}{2} \frac{\partial^2 G}{\partial V^2} h \right) dt + \frac{\partial G}{\partial V} h dW_t^{(1)}$$

Squaring and taking expectations, we note that all terms arising from the drift component are of order $dt^{3/2}$ or higher so that

$$\begin{aligned} \mathbf{E} [(dG_T(V_t))^2 | V_t] &= \left(\frac{\partial G}{\partial V} \Big|_{V_t} \right)^2 h^2 dt \\ &= \left(\frac{\frac{\partial}{\partial V} \mathbf{E} [Y_T^2 | V_t]}{2 \mathbf{E} [Y_T^2 | V_t]} \right)^2 h^2 dt \end{aligned}$$

From (3) we have

$$\mathbf{E} [Y_T^2 | V_t] = \int_t^T \exp \left(V_t e^{-gs} + \frac{h^2}{4g} (1 - e^{-2gs}) \right) ds$$

so that

$$\frac{\partial}{\partial V} \mathbf{E} [Y_T^2 | V_t] = \int_t^T e^{-gs} \exp \left(V_t e^{-gs} + \frac{h^2}{4g} (1 - e^{-2gs}) \right) ds$$

all of which can be computed by numerical integration.

To obtain the unconditional variance from the conditional one, it is only necessary to numerically integrate over the stationary distribution for V_t , namely $\mathbf{N} \left(0, \frac{h^2}{2g} \right)$.

$$\mathbf{E} [(dG_T(V_t))^2] = h^2 dt \int_{-\infty}^{\infty} \mathbf{E} [(dG_T(V_t))^2 | V_t] \phi \left(\frac{v}{\frac{h^2}{2g}} \right) dv$$

We also seek the correlation between log implied volatilities and the underlying, namely

$$\text{Corr} (dG_T(V_t), dY_t) = \frac{\mathbf{E} [dG_T(V_t) dY_t]}{\sqrt{\mathbf{E} [(dG_T(V_t))^2] \mathbf{E} [dY_t^2]}} \quad (18)$$

Applying Ito as above to $G_T(V_t)$ and dropping higher order terms, we have

$$\begin{aligned}
\mathbf{E} [dG_T(V_t)dY_t|V_t] &= \mathbf{E} \left[h \frac{\partial G}{\partial V} dW_t^{(1)} e^{\frac{V_t}{2}} dW_t | V_t \right] \\
&= h \rho dt \left(\frac{\partial G}{\partial V} \Big|_{V_t} \right) e^{\frac{V_t}{2}} \\
&= h \rho dt e^{\frac{V_t}{2}} \frac{\frac{\partial}{\partial V} \mathbf{E} [Y_T^2 | V_t]}{2 \mathbf{E} [Y_T^2 | V_t]}
\end{aligned}$$

which, once again, can be integrated over the stationary distribution of V_t to obtain the unconditional $\mathbf{E} [dG_T(V_t)dY_t]$.

Finally, we need

$$\begin{aligned}
\mathbf{E} [(dY_t)^2] &= dt \mathbf{E} \left[e^{\frac{V_t}{2}} \right] \\
&= dt e^{\frac{h^2}{4g}}
\end{aligned}$$

Note that all occurrences of dt in (18) cancel, and ρ appears only as a linear factor in the numerator.