Numerix LLC
Corporate Headquarters
99 Park Avenue
5th Floor
New York, NY 10016
Tel: +1 212 302 2220

Real-World Equity & Volatility Behavior: Implications for Economic Scenario Generation

December 2014

Written by **Andrew McClelland, PhD,** Financial Engineering, Numerix

In the risk-neutral world, all investments grow, on average, at the risk-free rate. However, in reality, risky positions earn a premium depending on their sensitivities to market factors. For instance, a call on an equity index is seen to earn a premium, while a put is seen to pay a premium.

Thus, while risk-neutral dynamics are ideal for producing valuations and hedge ratios, risk must be assessed with reference to real-world dynamics. Real-world dynamics incorporate risk premia, and this has immediate consequences for the generation of economic scenarios by insurance companies.

This paper examines the differences between risk-neutral dynamics and real-world dynamics, and the important role of risk premia. We highlight why real-world dynamics are necessary for risk analysis and scenario generation, and also explore the roles of the equity premium and volatility premium in stochastic volatility models of equity markets.

I. REAL-WORLD DYNAMICS VS. RISK-NEUTRAL DYNAMICS

Breaking it Down: ***Defining Risk-Neutral Dynamics***

From a technical perspective, how do we define risk-neutral dynamics? Risk-neutral dynamics are a "special" set of dynamics, which allow us to value instruments by discounting at the risk-free rate, *i.e.*

$$PV(t) = \tilde{E}[e^{-r(T-t)}\Psi(T)].$$

Here, $\tilde{E}[\cdot]$ instructs us to take the expectation with respect to the risk-neutral dynamics.

Risk-neutral dynamics for any traded security, say a stock S, must grow on average at the risk-free rate, *i.e.*

$$\frac{dS}{S} = rdt + \sigma d\widetilde{W}.$$

Importantly, we have an expected return of

$$\tilde{E}\left[\frac{dS}{S}\right] = rdt.$$

Defining Real-World Dynamics

A fundamental property of financial markets is that risky investments require compensation over the risk-free rate. Real-world dynamics of any security must reflect this, *i.e.*

$$\frac{dS}{S} = \mu dt + \sigma dW,$$

where μ represents the expected return on equity, including a risk premium, *i.e.* $(\mu - r) \geq 0$.

Using $E(\cdot)$ for real-world expectations, we now have:

$$E\left[\frac{dS}{S}\right] = \mu dt\,.$$

It is important to note that given an equity premium, even for simple equity processes, risk-neutral and real-world paths can diverge over long sample periods, as we will see momentarily. This clearly affects our assessment of risk under the two sets of dynamics, which leads us to question which set is appropriate.

For the purposes of evaluating risk, for example, we may want to know the probability of losing more than 10%, over a 1Y horizon. Which of the following probabilities would be most useful

$P(L > 10\%)$ or $\tilde{P}(L > 10\%)$?

Naturally, we would seek the "actual" probability of this loss taking place to assess our risk, rather than a hypothetical probability which is convenient for pricing purposes. Here, we should also recall that Historical VaR is implicitly a real-world risk metric, given that real-world dynamics govern historical market movements.

Let's take a look at the empirical evidence in the slides that follow to gain a clearer picture of the potential pitfalls of ignoring equity premium in our assessments.

*Diagram 1 – **Empirical Evidence on Equity Premium***

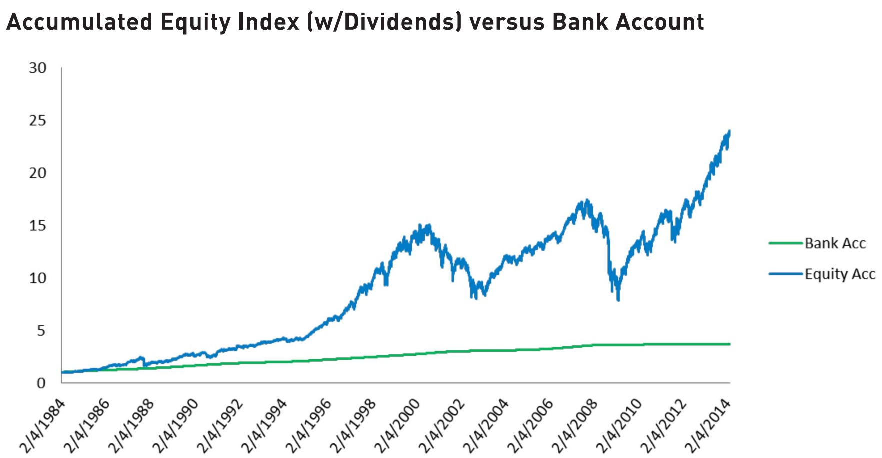

30 years of S&P500 accumulation (with dividend reinvestment) versus 30 years of rolling over at the risk-free (Fed Funds) rate (initialised at 1).

In Diagram 1, we have plotted an equity-linked account where we have taken a unit position in an S&P500 tracking fund and reinvested the dividends and an account where we have reinvested each day at the Fed Funds rate. Looking at the historical behaviour, we can see the effects of the equity premium on the growth of the investment, and note significant differences from the risk-free investment. But, what if we were to "trim" the historically-observed equity premium off the historical equity series? What would happen in our diagram?

Consider a simple estimate of the equity premium according to

$$\text{estimated eq. premium} = \frac{1}{T}\sum_{i=1}^{T}\frac{1}{\Delta_i}\left(\frac{\Delta S_i}{S_i} - r_i\Delta_i\right)$$

with

$$\frac{1}{\Delta_i}\left(\frac{\Delta S_i}{S_i} - r_i\Delta_i\right) \approx \frac{1}{dt}\left(\frac{dS(t)}{S(t)} - r(t)dt\right) = \mu(t) - r(t) + \cdots .$$

The historical average suggests $\mu - r \approx 5.89\%$ (on a continuously compounded basis).

*Diagram 3 – **Empirical Evidence on Equity Premium***

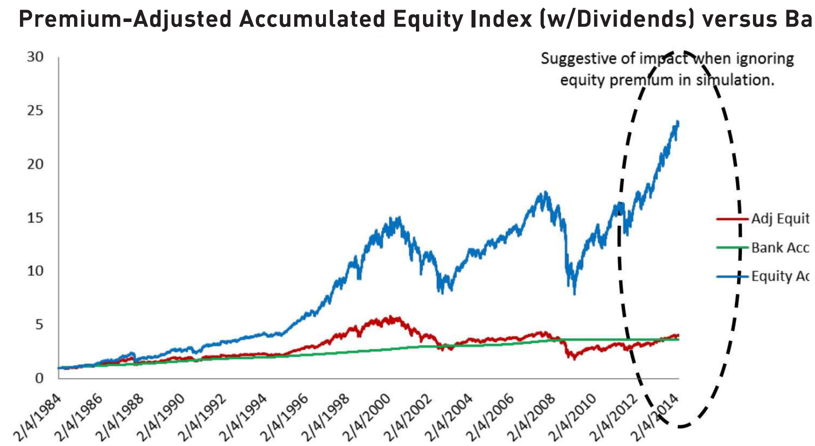

Adjusted equity accumulation index where the estimated equity premium has been trimmed of historical S&P500 returns.
Note the suggestive impact of ignoring equity premium in simulations, representing significant differences.

By removing the equity premium, we see that the equity-linked account, despite its larger fluctuations and greater risk, now has a similar total return to the rolling money-market account. Although this is to be expected given the use of an ex-post estimate of the equity premium, it highlights what could happen if we relied purely on risk-neutral scenarios when analyzing our risk. Indeed, the average of the simulated equity levels will be similar to that of the money-market account.

II. IMPLICATIONS FOR RISK PROFILES AND SCENARIOS

To further showcase this point, the following diagrams highlight what our economic scenarios would look like across our real-world and risk-neutral simulations—along with their differing risk profiles.

*Diagram 4 – **Simulated Real-World and Risk-Neutral Scenarios, Plus Risk Profiles***

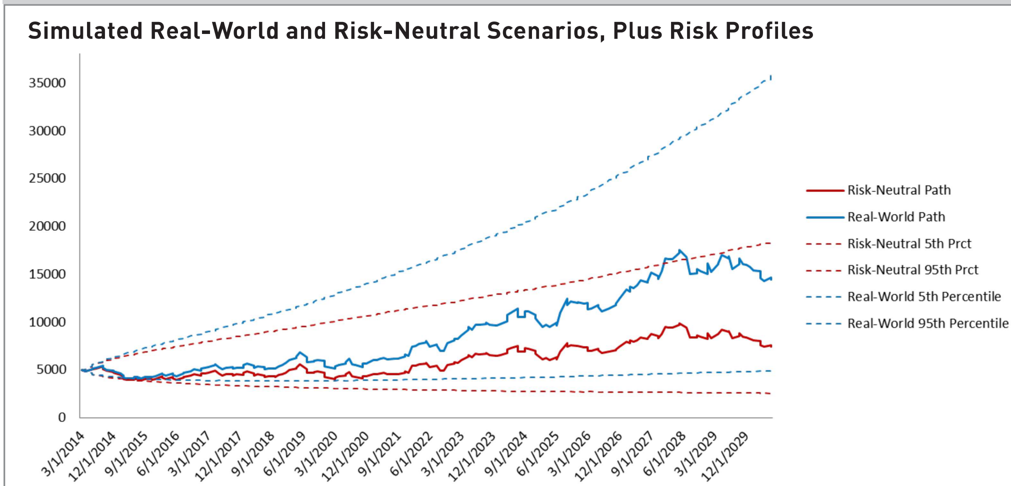

Simulated Black Scholes (1973) paths using both real-world and risk-neutral dynamics, along with percentiles of projected distributions. Monthly observations for 15 years. Percentiles were computed off samples of P=30,000 paths.

We can clearly observe that the real-world and risk–neutral distributions in Diagram 4 diverge quite dramatically. Therefore, again depending on which distribution is used, significantly different risk assessments will result.

III. WHAT DOES THIS MEAN FOR VARIABLE ANNUITIES AND STRUCTURED DERIVATIVES?

What would happen if we were using more complex derivative products? Let's have a look, for example, at risk for variable annuities and structured derivatives. For simple long positions in equity, the impact of ignoring the equity premium is somewhat straightforward. Since, variable annuity (VA) products behave like structured equity derivatives (given embedded guarantees), we can extend the simple risk premium analysis to derivatives to assess the impact on VAs. To keep matters simple, consider a European call, C, and a European put, P. We want to show that the expected return on the call, for example, is

$$E\left[\frac{dC}{C}\right] = (r + \text{loading on } dW \times \text{premium for } dW \text{ risk})dt.$$

We know that these satisfy, for instance, the famous Black Scholes PDE, *i.e.* for the call

$$rC = C_t + C_S rS + \left(\frac{1}{2}\right) C_{SS}\sigma^2 S^2,\ C(T) = (S-K)^+.$$

We can also apply Ito's Lemma to obtain the expected real-world dynamics for the derivative's value

$$dC = (C_t + C_S\mu S + (1/2)C_{SS}\sigma^2 S^2)dt + (C_S\sigma S)dW(t)$$

$$\Rightarrow \frac{dC}{C} = \left(\frac{C_t + C_S\mu S + (1/2)C_{SS}\sigma^2 S^2}{C}\right)dt + \left(\frac{C_S\sigma S}{C}\right)dW(t)$$

$$\Rightarrow E\left[\frac{dC}{C}\right] = \left(\frac{C_t + C_S\mu S + (1/2)C_{SS}\sigma^2 S^2}{C}\right)dt$$

^ **We need a way to relate μ *to* r!**

As we can see above, we will need a way to relate μ to r, as the terms highlighted in red are essentially the same, save for these parameters. To proceed, we again compare the real-world and risk-neutral dynamics we explored earlier, and we recall that a fundamental property of financial markets is that risky investments require compensation over the risk-free rate. Real-world dynamics of any security must reflect this as in the equation

$$\frac{dS}{S} = \mu dt + \sigma dW = (r + \theta\sigma)dt + \sigma dW.$$

Here μ represents the expected return on equity, including a risk premium, *i.e.* $(\mu - r) = \theta\sigma$ (with θ representing the "price per unit of equity risk").

Using $E(\cdot)$ for real-world expectations, we now have

$$E\left[\frac{dS}{S}\right] = \mu dt = (r + \theta\sigma)dt,$$

where we note that θ is essentially the well-known Sharpe ratio

$$\frac{\mu - r}{\sigma} = \frac{\theta\sigma}{\sigma} = \theta.$$

Returning to our expected return equation for the call and substituting in for yields

$$E\left[\frac{dC}{C}\right] = \left(\frac{C_t + C_S(r+\theta\sigma)S + (1/2)C_{SS}\sigma^2 S^2}{C}\right)dt,$$

and after taking advantage of the Black Scholes PDE we obtain

$$E\left[\frac{dC}{C}\right] = \left(\frac{rC + \theta\sigma C_S S}{C}\right)dt,$$

which of course simplifies to

$$E\left[\frac{dC}{C}\right] = \left(r + \left(\frac{C_S\sigma S}{C}\right)\theta\right)dt.$$

If we note that the factor

$$\frac{C_S\sigma S}{C}$$

is in fact the return volatility for the call, which we saw earlier after applying Ito's Lemma, it is clear then that the premium on the call is the product of its loading onto raw equity risk multiplied by the premium for raw equity risk, i.e.

$$E\left[\frac{dC}{C}\right] = (r + \text{loading on } dW \times \text{premium for } dW \text{ risk})dt$$

Importantly, calls have positive Deltas, as they are "long equity"

$$C_S > 0 \Rightarrow \frac{C_S\sigma S}{C} > 0.$$

Hence, calls "earn" a premium because of their positive loading. The same arguments can be made for put contracts, but

$$P_S < 0 \Rightarrow \frac{P_S\sigma S}{P} < 0.$$

Thus, puts "pay" a premium (like an insurance contract which they indeed are).

The examples that follow visually highlight the real-world and risk-neutral risk profile differences for call and put options.

*Diagram 5 – **Real-world vs. Risk Neutral Profiles for a Call Option***

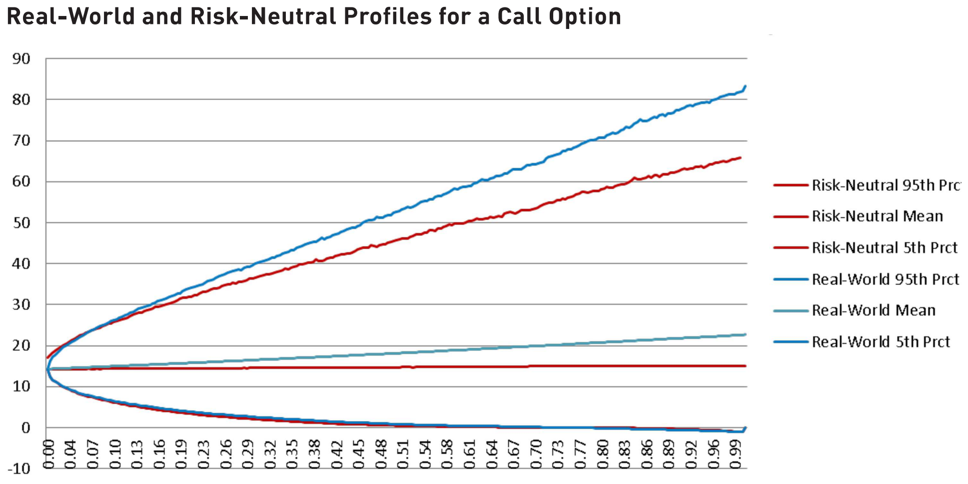

Real-world and risk-neutral risk profiles for a call option according to the Black Scholes (1973) dynamics.
Here $S_0 = 100$, $K = 100$, $\mu = 0.15$, $r = 0.05$ and $\sigma = 0.3$.

*Diagram 6 – **Real-world vs. Risk Neutral Profiles for a Put Option***

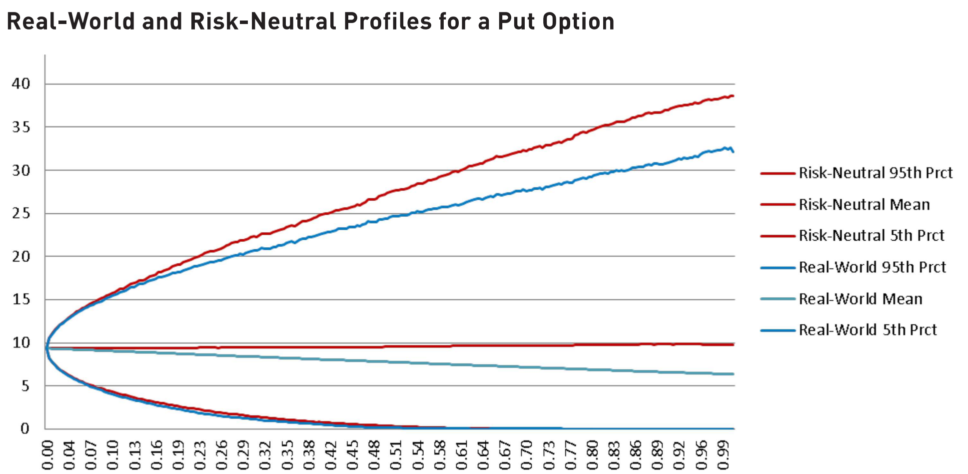

Real-world and risk-neutral risk profiles for a put option according to the Black Scholes (1973) dynamics.
Here $S_0 = 100$, $K = 100$, $\mu = 0.15$, $r = 0.05$ and $\sigma = 0.3$.

We can clearly observe that the risk profile tilts upwards for calls to reflect the positive premium they earn, and tilts downwards for puts to reflect the negative premium that they earn (pay). We should also note that the previous examples are for relatively short risk horizons of only one year, and that the differences typically become larger as the time horizon grows.

We have so far worked only with simple products and positions, but the differences between physical and risk-neutral behavior can become quite a lot more significant, depending on the type of product, and the product's particular sensitivities. It can also become highly pronounced when leverage enters the picture!

For example, let's look at a case in which we are buying 10 units of the underlying call, but using 90% leverage (*i.e.* only putting up 10% of the capital for each call, and thus having the same initial capital outlay). In Diagram 7 we see that the difference between the upper percentiles for the leveraged position is now in the hundreds of dollars, as opposed to tens of dollars, for a position with the same amount of up-front capital.

*Diagram 7 – **The Impact of Leverage on Real-World & Risk-Neutral Risk Profiles***

Real-World and Risk-Neutral Profiles for a Call Option with 90% Leverage

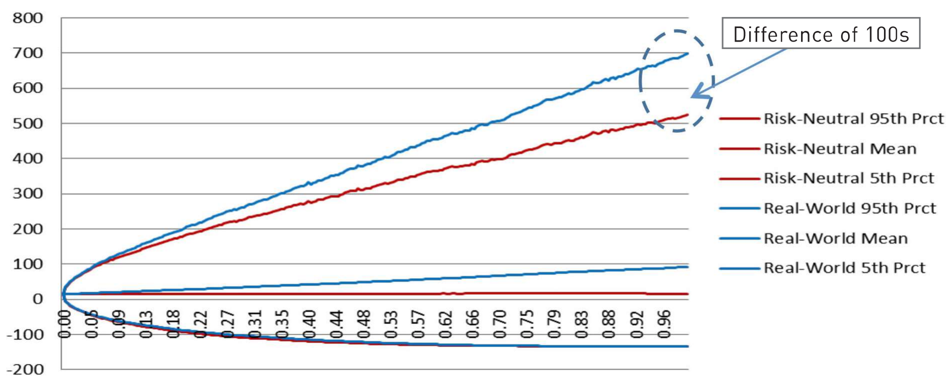

Significantly different real-world and risk-neutral risk profiles for a leveraged position (90%) in call options, according to the Black Scholes (1973) dynamics.
Here $S_0 = 100$, $K = 100$, and $\sigma = 0.3$.

IV. STOCHASTIC VOLATILITY AND THE VOLATILITY PREMIUM

Up until this point, our study has considered only a simple one-factor equity model. Things get a little more complex in multi-factor models, in which each factor attracts a premium. Arguably, the most important additional factor for equity markets is stochastic volatility—as it is an all-pervasive feature of the financial markets. Therefore, let's start our study by taking a look at a diagram of 10-day rolling volatility of the S&P500 over the past 30 years. It is easy to observe periods of volatility clustering and a tendency to revert to a long-run mean.

*Diagram 8 – **Bringing Volatility Into the Picture***

10-Day Rolling S&P500 Volatilities (Annualised)

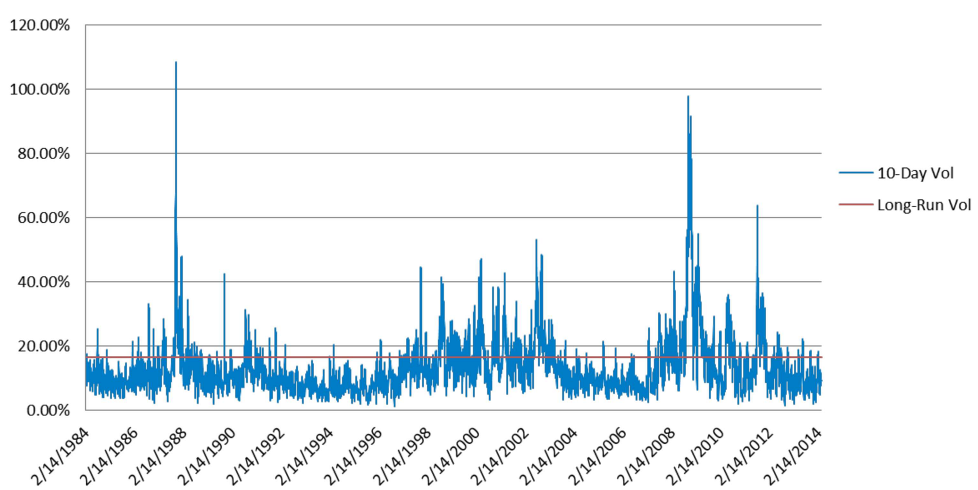

10-day rolling volatility of the S&P500 over the past 30 years. The 10-day volatilities have been quoted on an annualised basis.

When stochastic volatility is present, derivatives also earn/pay a premium for their sensitivities to volatility. Take for instance the risk-neutral Heston (1993) model

$$\frac{dS}{S} = rdt + \sqrt{V}d\widetilde{W}_1$$

$$dV = \tilde{\kappa}(\tilde{V}_\infty - V)dt + \xi\sqrt{V}d\widetilde{W}_2$$

with $< d\widetilde{W}_1, d\widetilde{W}_2 >= \rho < 0.$

The Heston model is popular given its ability to produce both the heavy-tailed return distributions that are commonly observed in financial markets, and also the volatility persistence discussed earlier. These can easily be demonstrated via simulation, as the following graphs suggest.

*Diagram 9 – **Looking at Equity and Volatility Distributions Side-by-Side***

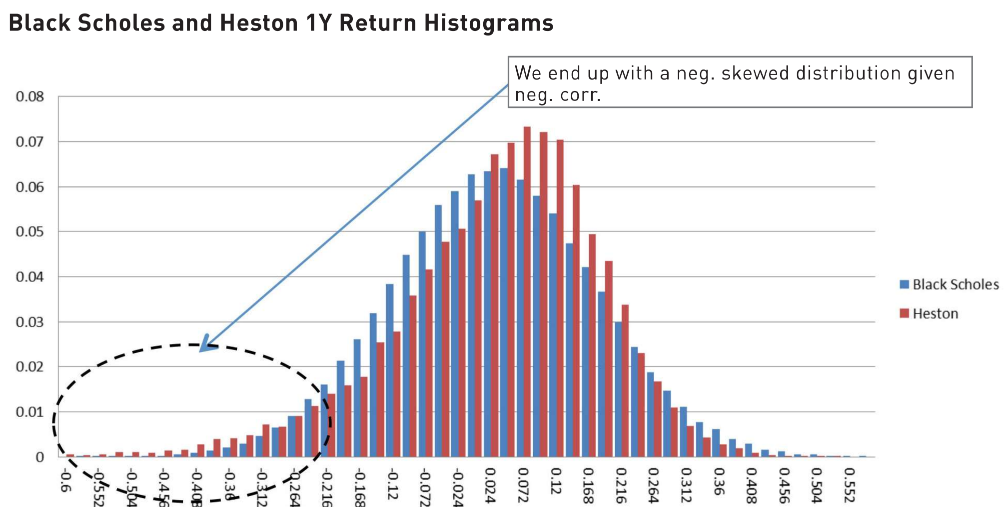

A comparison of the Black Scholes (1973) and Heston (1993) 1-year log-return distributions. Using P=15,000 simulations.

*Diagram 10 – **True Heston and Rolling Variance***

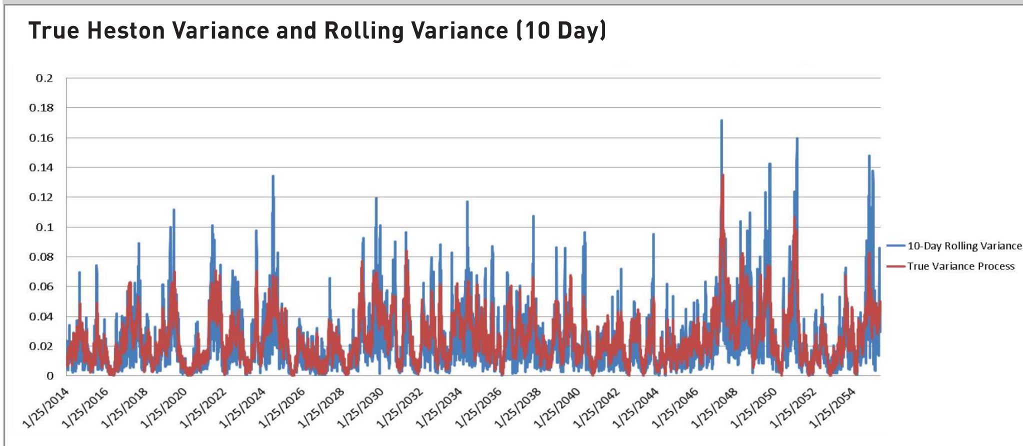

A simulation of volatilities and rolling 10-day volatilities using the Heston (1993) model.

Let's review the *risk-neutra*l Heston model, as compared with the *real-world* Heston model directly below it. First the risk-neutral dynamics,

$$\frac{dS}{S} = r dt + \sqrt{V} dW_1$$

$$dV = \tilde{\kappa}\left(\tilde{V}_\infty - V\right)dt + \xi\sqrt{V} dW_2$$

with $< dW_1, dW_2 >= \rho$.

Now the real-world dynamics model:

$$\frac{dS}{S} = (r + \phi V)dt + \sqrt{V}dW_1$$

$$dV = \kappa(V_\infty - V)dt + \xi\sqrt{V}dW_2$$

with $< dW_1, dW_2 >= \rho$, and now $\tilde{\kappa} = \kappa - \gamma$, $\frac{\kappa}{\tilde{\kappa}} = \frac{\tilde{V}_\infty}{V_\infty}$.

Here, $\phi > 0$ and $\gamma > 0$ govern equity and volatility premium, respectively. Like the equity premium reflects the fact that investors (on average) are adverse to negative movements in equity prices, the volatility premium reflects that they are also adverse to positive movements in volatility, and are willing to pay a premium to hedge against such moves! This must be reflected in derivative prices (which are vol hedges), and this is facilitated by warping the risk-neutral volatility dynamics.

To offer some empirical evidence supporting the notion that the volatility premium is indeed present, consider the VIX and 22-day ahead S&P500 return variance. The VIX is defined as the risk-neutral expected 22-day return variance

$$\text{VIX} = \tilde{E}[\text{22 day return variance}].$$

If there is no volatility premium, we have

$$\underline{\tilde{E}[\text{22 day return variance}] = E[\text{22 day return variance}]}$$

no vol premium $\Longrightarrow$ *same* vol dynamics

Thus, the VIX should be an *unbiased* predictor of future S&P500 return variance, and we can assess whether this is the case graphically.

*Diagram 11 – **VIX as predictor of future S&P500 return variance***

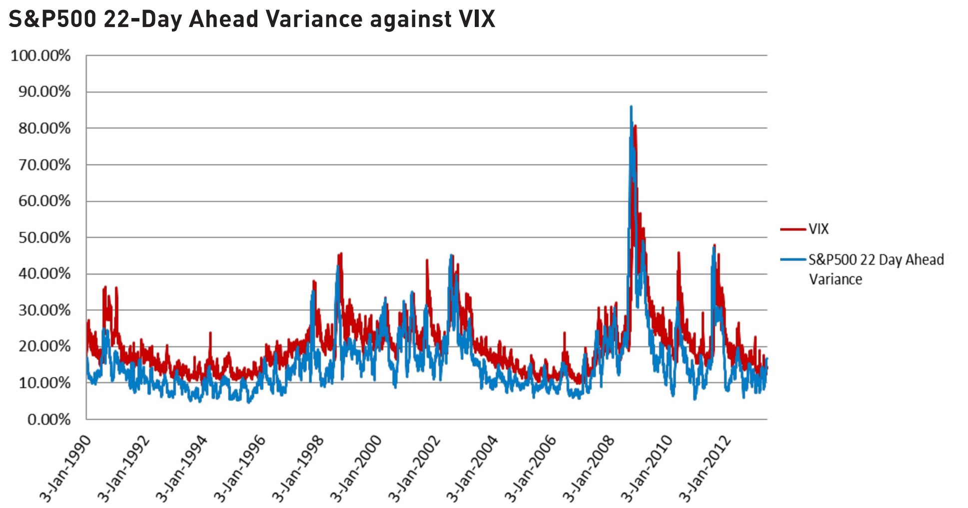

CBOE VIX index against 22-day ahead variance of S&P500 returns over the period Jan 1990-Jun 2013.

Figures are reported as annualised volatilities.

Recalling our choice of parameter constraints for the governing equity and volatility premia, $\phi > 0$ and $\gamma > 0$, we have

$$\tilde{\kappa} = \kappa - \gamma < \kappa \Rightarrow \tilde{\kappa} < \kappa$$

and

$$\frac{\kappa}{\tilde{\kappa}} > 1 \Rightarrow \frac{\tilde{V}_\infty}{V_\infty} > 1.$$

As such, we have lower reversion and higher long-run volatility under the risk-neutral dynamics!

Skipping the math and setting $\rho = 0$ for simplicity, we can arrive at a premium breakdown of the call premium:

$$\Rightarrow \boldsymbol{E\left[\frac{dC}{C}\right] = \left(r + \left(\frac{C_S\sqrt{V}S}{C}\right)\phi\sqrt{V} - \left(\frac{C_V\xi\sqrt{V}}{C}\right)\frac{\gamma\sqrt{V}}{\xi}\right)dt}$$

$\boxed{\phi \text{ controls equity component}}$ $\boxed{\gamma \text{ controls volatility component}}$

We have essentially the same breakdown as for the single-factor Black Scholes case; there is a sensitivity to a factor, multiplied by the premium earned for taking exposure to that factor. Importantly, the vegas for both calls and puts are positive, *i.e.* $C_V > 0$ and $P_V > 0$, and so both are "vol hedges".

To see the impact explicitly, consider a Delta-hedged position in a call:

$$E\left[\frac{dC^\Delta}{C^\Delta}\right] = \left(\underbrace{r + \left(\frac{C_S^\Delta\sqrt{V}S}{C^\Delta}\right)\phi\sqrt{V}}_{\boxed{= 0 \text{ as } C_S^\Delta = 0}} - \left(\frac{C_V^\Delta\xi\sqrt{V}}{C^\Delta}\right)\frac{\gamma\sqrt{V}}{\xi}\right)dt$$

$$E\left[\frac{dC^\Delta}{C^\Delta}\right] = \left(r - \left(\frac{C_V^\Delta\xi\sqrt{V}}{C^\Delta}\right)\frac{\gamma\sqrt{V}}{\xi}\right)dt.$$

This is a "pure" vol hedge, and thus must pay a premium. In the diagrams that follow, we visually highlight the risk profiles of real-world and risk-neutral delta-hedged call and put options to note the differences. The delta-hedged call option clearly grows less quickly under the physical measure where the volatility premium is accounted for. As the put option is also long vega, it exhibits similar behavior.

*Diagram 12 – **Risk Profiles of Delta-Hedged Call Option***

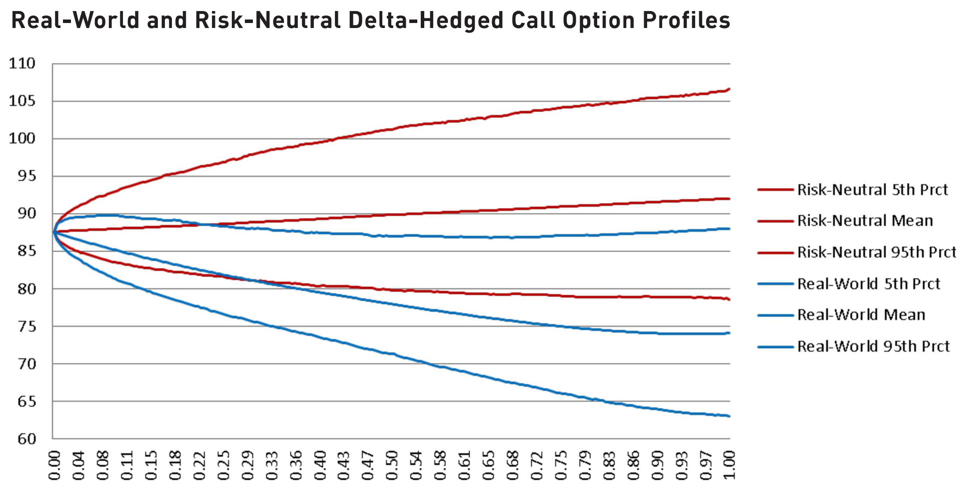

Real-world and risk-neutral risk profiles for a delta-hedged put option according to the Heston (1993) dynamics.

CONCLUSION

Risk-neutral dynamics are absolutely the correct dynamics to use for valuations and hedging. However, we have clearly observed in this study why it is so important to analyze risk measures and scenarios through the real-world lens. For derivatives and related structures, real-world profiles can substantially deviate from risk-neutral profiles—depending on such factors as product, sensitivities and leverage—leading to inaccurate risk profile assessments. Real-world dynamics incorporate risk premia, and as we have seen in this study, this has immediate consequences for the generation of economic scenarios by insurance companies.

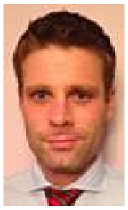

AUTHOR BIOGRAPHY

Andrew McClelland, PhD, Financial Engineering, Numerix

Andrew McClelland's work at Numerix focuses on counterparty credit risk issues including valuation adjustments and counterparty exposure production for structured products, and numerical methods for efficient production of risk profiles under the real-world measure. Dr. McClelland earned his PhD in finance at the Queensland University of Technology for a thesis on financial econometrics. He considered markets exhibiting crash feedback, option pricing for such markets, and parameter estimation for such markets using particle filtering methods. His work has been published in the Journal of Banking and Finance and the Journal of Econometrics.

Ready to elevate your business?

Get started →

More than 550 clients and 90 partners across 52 countries rely on Numerix analytics. Headquartered in the financial capital of the world, with offices around the globe, Numerix brings together unparalleled expertise combined with local market understanding.

numerix

Numerix LLC Corporate Headquarters

100 Park Avenue
15th Floor
New York, NY 10017
Tel: +1.212.302.2220
Fax: +1.212.302.6934
www.numerix.com