

PRICING AND HEDGING VARIABLE-VOLUME SWAPS

HEMANTH HARIHARAN

1. INTRODUCTION AND REVIEW

The energy component of Full Requirements (FR) deals is typically a short position in a variable-volume physical swap, that is valued by calculating a shaped energy price and a load following adder (LFA). The LFA is one of the most nuanced and mathematical concepts involved in pricing a deal on the FR desk. It is sometimes interpreted as something we charge to cover the mismatch between the flat hedges we procure and the realized load, but I believe that is an incorrect interpretation. I have come across an interpretation where the LFA is equivalent to the value of an option to follow the load, and we will indeed prove to ourselves after this derivation that the LFA is the measure of the optionality embedded into a variable-volume swap. After having spent some time reviewing the approved methodology paper by risk, I am expanding and rewriting the derivation in more detail, in the hope that analysts who are mathematically inclined and curious about the theory behind the valuation of FR deals, benefit from it.

1.1. Review. Before delving into the derivation, basic concepts from probability, statistics, linear algebra and stochastic calculus need to be revisited including the following:

- (1) Normal Distribution
- (2) Lognormal Distribution
- (3) Wiener Processes and Geometric Brownian Motion
- (4) Ito's Lemma
- (5) Cholesky Decomposition

2. VALUATION

2.1. Risk-neutral valuation. The Risk Analytics (RA) side of the FR desk prices deals using a risk-neutral valuation approach where the expected value of the future payoff of the derivative is zero. For the sake of simplicity, we can assume that the deal is energy-only with no capacity, ancillary services or RECs. We can also start off by assuming that this is a "1-hour" deal, and then generalize this to a real-world case towards the end of the derivation.

Let's define the following variables

- (1) L - Load (Volume)
- (2) P - Floating Price
- (3) k - Fixed Contract Price

In a typical FR deal, we sell L MWhs of power at the fixed contract price k while we buy power at the floating market P price. The payoff of this "1-hour" deal is $(k - P)L$

Let's say we want to find the fair contract price k such that $\mathbb{E}[(k - P)L]$ is 0. Expanding the term inside the expectation, we get $\mathbb{E}[(k.L) - (P.L)] = 0$. Note that while P and L are random variables, k is deterministic.

Hence we get $k = \frac{\mathbb{E}[P.L]}{\mathbb{E}[L]}$. It has been statistically observed by the team that P and L tend to be correlated random variables in practice. Hence, by using the definition of covariance, we get

$$k = \mathbb{E}[P] + \frac{\text{cov}(P, L)}{\mathbb{E}[L]} \approx \text{SEP} + \text{LFA} \quad (2.1)$$

When we extrapolate this oversimplified "1-hour" deal to a more general case, the first term is known as the shaped-energy price (SEP), which simply put, is the load-weighted average price of energy. The second term, is the LFA, which is the main topic of discussion for this paper. Calculating the covariance of price and load would require us to make assumptions on the nature of these variables which we shall do in the subsequent section.

2.2. Geometric Brownian Motion. The risk-approved methodology to calculate the LFA assumes that price and load follow a correlated geometric Brownian motion. While this is a fairly standard model for stock prices, applying this to electricity prices and load is something worth researching further. Nevertheless, it is something the team has gotten comfortable with over the years, and will form the basis for this paper as well.

Using the definitions of geometric Brownian motion, we get

$$\frac{dP}{P} = \mu_P dt + \sigma_P dz_1 \quad (2.2)$$

$$\frac{dL}{L} = \mu_L dt + \sigma_L dz_2 \quad (2.3)$$

where μ_L and μ_P are the means of the return on load and price, while σ_L and σ_P are the standard deviations of the return on load and price respectively. dz_1 and dz_2 are Wiener processes with a correlation of ρ . It is important to note that load and price volatility are, by definition, the standard deviations of the return on load and price, and not load and price itself. The discretized versions of (2.2) and (2.3) give us

$$\frac{\delta P}{P} = \mu_P \delta t + \sigma_P \sqrt{\delta t} \phi_1 \quad (2.4)$$

$$\frac{\delta L}{L} = \mu_L \delta t + \sigma_L \sqrt{\delta t} \phi_2 \quad (2.5)$$

where ϕ_1 and ϕ_2 are standard normal variables with a correlation of ρ . Note that over a period δt , the mean of the return on (say) load is $\mu_L \delta t$ while the variance of the return on load is $\sigma_L^2 \delta t$. This arises as a consequence of the Markov property of geometric Brownian motion, where the future state of the process is only dependent on the current state and independent of history. Since the variance per unit time is σ_L^2 , the variance over a period δt is $\sigma_L^2 \delta t$ (since the variance of the sum of independent random normal variables is the sum of the variances of the variables) and consequently, the standard deviation over a period δt is $\sigma_L \sqrt{\delta t}$.

Applying Ito's Lemma to (2.2) and (2.3) leads us to the lognormal property of load and prices where

$$\log \frac{P_t}{P_o} \sim N \left(\left(\mu_P - \frac{\sigma_P^2}{2} \right) t, \sigma_P \sqrt{t} \right) \quad (2.6)$$

$$\log \frac{L_t}{L_o} \sim N \left(\left(\mu_L - \frac{\sigma_L^2}{2} \right) t, \sigma_L \sqrt{t} \right) \quad (2.7)$$

Although a complete review of the theoretical underpinnings of geometric Brownian motion is beyond the scope of this document, the above equations should be sufficient to help us derive the equation for the LFA. For more depth in understanding, I highly recommend reviewing the chapters on Wiener processes and stock price modeling in [1].

2.3. Cholesky Decomposition. The Cholesky decomposition method allows us to generate correlated random variables. Given a vector of independent random variables $X = [\epsilon_1, \epsilon_2]^T$, Cholesky decomposition creates a new vector Z of correlated random variables as follows,

$$Z = LX \quad (2.8)$$

and

$$L = \begin{bmatrix} 1 & 0 \\ \rho & \sqrt{1 - \rho^2} \end{bmatrix} \quad (2.9)$$

which gives us

$$Z = \begin{bmatrix} \epsilon_1 \\ \rho \epsilon_1 + \sqrt{1 - \rho^2} \epsilon_2 \end{bmatrix} \quad (2.10)$$

2.4. Calculating the covariance. Now that we've established the foundation for the derivation, we can go ahead and express P and L as functions of time, and calculate their covariance and ultimately the LFA. Expressing load and price as a function of time using (2.6), (2.7) and applying Cholesky decomposition using (2.10), we get

$$P_t = P_0 \exp \left(\left(\mu_p - \frac{\sigma_P^2}{2} \right) t + \sigma_P \sqrt{t} \epsilon_1 \right) \quad (2.11)$$

and

$$L_t = L_0 \exp \left(\left(\mu_L - \frac{\sigma_L^2}{2} \right) t + \sigma_L \sqrt{t} (\rho \epsilon_1 + \sqrt{1 - \rho^2} \epsilon_2) \right) \quad (2.12)$$

Note that

$$E[P_t] = P_0 \exp(\mu_P t) \quad (2.13)$$

and

$$E[L_t] = L_0 \exp(\mu_L t) \quad (2.14)$$

The covariance of P and L is now calculated, by definition, as

$$\text{cov}(P_t, L_t) = E[P_t L_t] - E[P_t] E[L_t] \quad (2.15)$$

Since P and L are both functions of time, the covariance will also be a function of time. Substituting (2.11) and (2.12) into $E[P_t L_t]$ in (2.15), by making use of the fact that ϵ_1 and ϵ_2 are independent random variables and grouping terms, we get

$$\begin{aligned} E[P_t L_t] = P_0 L_0 \exp((\mu_p + \mu_L)t) \exp \left(\frac{-\sigma_P^2 t - \sigma_L^2 t}{2} \right) & E \left[\exp((\sigma_P \sqrt{t} + \rho \sigma_L \sqrt{t}) \epsilon_1) \right] \\ & E \left[\exp(\sqrt{1 - \rho^2} \sigma_L \sqrt{t} \epsilon_2) \right] \end{aligned} \quad (2.16)$$

Since ϵ_1 and ϵ_2 are standard normal random variables, the terms inside the exponentials in (2.16) are normal random variables, which also implies that the exponentials of these are lognormal variables, by definition. Recalling that the expectation of a lognormal distribution with parameters μ and σ is $\exp(\mu + \frac{\sigma^2}{2})$, we get

$$E[P_t L_t] = P_0 L_0 \exp((\mu_p + \mu_L)t) \exp\left(\left(\frac{-\sigma_P^2}{2} + \frac{-\sigma_L^2}{2}\right)t\right) \exp\left(\frac{(\sigma_P \sqrt{t} + \rho \sigma_L \sqrt{t})^2}{2}\right) \exp\left(\frac{(\sqrt{1 - \rho^2} \sigma_L \sqrt{t})^2}{2}\right) \quad (2.17)$$

Expanding and simplifying (2.17), we get

$$E[P_t L_t] = P_0 L_0 \exp((\mu_p + \mu_L)t) \exp(\rho \sigma_P \sigma_L t) \quad (2.18)$$

Substituting (2.18), (2.13) and (2.14) in (2.15), we get

$$\text{cov}(P_t, L_t) = P_0 L_0 \exp((\mu_p + \mu_L)t) (\exp(\rho \sigma_P \sigma_L t) - 1) = E[P_t] E[L_t] (\exp(\rho \sigma_P \sigma_L t) - 1) \quad (2.19)$$

Substituting (2.19) in (2.1), we get

$$LFA = \frac{\text{cov}(P_t, L_t)}{E[L_t]} = \frac{E[P_t] E[L_t] (\exp(\rho \sigma_P \sigma_L t) - 1)}{E[L_t]} = E[P_t] (\exp(\rho \sigma_P \sigma_L t) - 1) \quad (2.20)$$

Note that the correlation and the load volatility is something we calculate internally, while we get the expected value of the future price from the futures/forwards desk and the price volatility from the options desk. I have also empirically verified the covariance in [3] by simulating correlated prices and volumes to confirm (2.20).

2.5. Load Volatility. On this desk, we calculate load volatility by the approach described below:

We first calculate the variance of load, $\text{var}(L_t)$, and the expected load based on our forecast, $E[L_t]$, and then use these variables to calculate the load volatility, which as mentioned earlier, is the standard deviation of the return on load. Using (2.7) and the properties of a lognormal distribution, we get

$$\text{var}[L_t] = (E[L_t])^2 (\exp(\sigma_L^2 t) - 1) \quad (2.21)$$

Solving for σ_L in (2.21), we get

$$\sigma_L = \sqrt{\frac{1}{t} \log \left(\frac{\text{var}[L_t]}{(E[L_t])^2} + 1 \right)} \quad (2.22)$$

This is the equation used to compute the load volatility from the standard deviation and expectation of load. Note that in practice, $\frac{\text{var}[L_t]}{(E[L_t])^2}$ tends to be lesser than 1. If we use the Taylor series expansion about 0 of $\log(1+x) \approx x$ by ignoring the higher order terms, we can simplify (2.22) as

$$\sigma_L \approx \frac{1}{\sqrt{t}} \frac{\text{stdev}(L_t)}{E[L_t]} \quad (2.23)$$

Note that I am explicitly stating the standard deviation of L_t and not using the symbol σ_L as the latter is the standard deviation of the return on load and not load itself. With all of these variables calculated, one could plug them into (2.1) and calculate the contract price that we should charge, such that on average, we make no money on the deal. The term that the team uses for this is the P-50 price. Strictly speaking, the P-50 value of a distribution is its median, but the team uses the term interchangeably with the mean.

2.6. Extrapolating to a real deal. Congratulations for following along thus far! We have already covered the bulk of the derivation, and all that remains is to extrapolate this to a real-world deal.

Let us say we first want to find the fair contract price CP for a deal that runs for M months and that we split each month into B blocks (say WD1, WD2, WD3, WD4, WE1... for example). The payoff from this deal is $\sum_{m=1}^M \sum_{b=1}^B (k - P_{b,m})L_{b,m}$. By the principle of risk-neutral valuation, setting the expected payoff to 0, we get

$$E \left[\sum_{m=1}^M \sum_{b=1}^B (k - P_{b,m})L_{b,m} \right] = 0 \quad (2.24)$$

Rearranging the terms and by applying linearity of expectation to (2.24), we can solve for the contract price CP as follows:

$$k = \frac{\sum_{m=1}^M \sum_{b=1}^B E [P_{b,m}L_{b,m}]}{\sum_{m=1}^M \sum_{b=1}^B E [L_{b,m}]} \quad (2.25)$$

Using the definition of covariance, (2.25) can be rewritten as:

$$k = \frac{\sum_{m=1}^M \sum_{b=1}^B E[P_{b,m}]E[L_{b,m}]}{\sum_{m=1}^M \sum_{b=1}^B E[L_{b,m}]} + \frac{\sum_{m=1}^M \sum_{b=1}^B cov(P_{b,m}, L_{b,m})}{\sum_{m=1}^M \sum_{b=1}^B E[L_{b,m}]} = SEP + LFA \quad (2.26)$$

As before, the first term is the SEP and the second term is the LFA. Note that this time though, to calculate the SEP, we would need to shape the forward price for each month into each block and forecast the load for each block of each month. To calculate the LFA, we can substitute (2.19) into (2.26) as follows:

$$LFA = \frac{\sum_{m=1}^M \sum_{b=1}^B cov(P_{b,m}, L_{b,m})}{\sum_{m=1}^M \sum_{b=1}^B E[L_{b,m}]} = \frac{\sum_{m=1}^M \sum_{b=1}^B E[P_{b,m}]E[L_{b,m}](exp(\rho_{b,m}(\sigma_P)_{b,m}(\sigma_L)_{b,m}t_m) - 1)}{\sum_{m=1}^M \sum_{b=1}^B E[L_{b,m}]} \quad (2.27)$$

We're almost there! The only piece of information we have glossed over is the time value of money. Strictly speaking, the present value of the expected payoff in (2.24) should be set to 0, by discounting the cash flows at the appropriate rates. This leads to a minor modification to (2.27) and the final formula for LFA as follows:

$$LFA = \frac{\sum_{m=1}^M \sum_{b=1}^B D_m E[P_{b,m}]E[L_{b,m}](exp(\rho_{b,m}(\sigma_P)_{b,m}(\sigma_L)_{b,m}t_m) - 1)}{\sum_{m=1}^M \sum_{b=1}^B D_m E[L_{b,m}]} \quad (2.28)$$

where D_m is the appropriate discounting factor. On our desk, we assume continuous compounding which leads to a form $\exp(-rt)$, where r is the risk-free interest rate plus a spread based on our credit rating. This concludes the derivation of the formula for the LFA. It is important to note that we only get ON-peak hub volatilities from the options desk, and use price volatility multipliers to transform these into monthly zonal volatilities for the different time blocks.

2.7. Price-Load Correlation. The term ρ in (2.9) is the correlation between ϕ_1 and ϕ_2 in (2.4) and (2.5). However, we need to come up with a way to express the correlation as a function of load and price in order to be able to calculate it statistically. Since ϕ_1 and ϕ_2 are standard normal variables, their correlation equals their covariance. Using this fact, rearranging and solving for ϕ_1 and ϕ_2 in (2.4) and (2.5), we get

$$\rho = cov(\phi_1, \phi_2) = cov\left(\frac{\log \frac{P_t}{P_0} - \mu_P t + \frac{\sigma_P^2}{2} t}{\sigma_P \sqrt{t}}, \frac{\log \frac{L_t}{L_0} - \mu_L t + \frac{\sigma_L^2}{2} t}{\sigma_L \sqrt{t}}\right) \quad (2.29)$$

Simplifying this, we observe that the constants drop off and we end up with $\rho = \text{correl}(\log P_t, \log L_t)$. which explains why the team currently estimates correlations between the natural logarithm of load and price.

2.8. Types of options and volatility. There are two main types of options that trade on an underlying of power - daily and monthly options. Typically, the options desk uses the Black-Scholes equation or some variant of it, to back-calculate the implied daily and monthly volatility from the premiums the respective options are trading at.

The payoff of the daily options is very similar to that of an Asian option (although ICE refers to them as lookback options for some reason, that does not match with the textbook definition of lookback options), with the option paying off whenever the average Day-Ahead (DA) LMP exceeds the strike price (in the case of a call option). The volatility implied by daily options hence corresponds to that of the DA spot market.

While the daily options have an underlying of the spot market, the monthly options are options on futures/swaps. Hence, the volatility implied by monthly options corresponds to that of the futures.

We will now derive an equation relating the daily and monthly volatility.

Let's define the following variables

- (1) σ_D - Daily volatility (implied by daily options)
- (2) σ_M - Monthly volatility (implied by monthly options)
- (3) σ_C - Cash volatility (intra-month spot volatility)
- (4) T - Time from today (or the evaluation date) to the middle of the month, for which we want to calculate the LFA
- (5) T_{start} - Time from today to the start of the month in question
- (6) F_0 - Forward price today for time T
- (7) F_{end} - Forward price at the end of the month before the month in question
- (8) P_{start} - Spot price at the start of the month in question
- (9) P_T - Spot price at the middle of the month in question

$$\frac{P_T}{F_0} = \frac{P_{start}}{F_0} \frac{P_T}{P_{start}} \approx \frac{F_{end}}{F_0} \frac{P_T}{P_{start}} \quad (2.30)$$

The above approximation can be made since, as the forward contract gets closer to maturity and rolls into the spot market, the forward price will converge to the spot price. Taking the natural logarithm of both sides of (2.30), we get

$$\ln \frac{P_T}{F_0} \approx \ln \frac{F_{end}}{F_0} + \ln \frac{P_T}{P_{start}} \quad (2.31)$$

We then take the variance of both sides of (2.31) to get

$$\text{var}[\ln \frac{P_T}{F_0}] \approx \text{var}[\ln \frac{F_{end}}{F_0}] + \text{var}[\ln \frac{P_T}{P_{start}}] \quad (2.32)$$

Note that in the above step, we have used the Markov property of geometric Brownian motion to assume that the terms on the right-hand-side of (2.31) are independent. From (2.6), we see that the variance of the log-return of price is the square of the volatility multiplied by time. Applying this result to (2.32) and using the appropriate volatility, we get

$$\sigma_D^2 T = \sigma_M^2 T_{start} + \sigma_C^2 (T - T_{start}) \quad (2.33)$$

Given the daily and monthly volatility implied by the options, one can use (2.33) to calculate the cash volatility σ_C , which can be used to calculate the LFA purely based on the intra-month volatility of the spot market. Note that the daily volatility is a blend of the volatility of forward and the spot market, while the monthly and cash volatility are exclusive to the forward and spot market respectively.

2.9. Pricing energy-pass through deals. On the FR desk, some deals are energy-pass through where we're paid the Real-Time (RT) price. However, we choose to schedule load on the DA market, which exposes us to the DART spread. As before, let us value a simplified "1-hour" energy pass-through deal. We would be buying L MWhs at the DA price and selling it at the RT price. We want to find the fair price of the deal k such that the expected value of $(k + RT - DA)L$ equals 0. Solving for k , we get

$$k = \frac{E[(DA - RT)L]}{E[L]} = \frac{E[DA.L] - E[RT.L]}{E[L]} \quad (2.34)$$

From (2.13), (2.14) and (2.18), we know that $E[PL] = E[P]E[L]exp(\rho\sigma_P\sigma_L t)$. Applying this to (2.34), we get

$$k = \frac{E[DA]E[L]exp(\rho_{DA}\sigma_{DA}\sigma_L t)}{E[L]} - \frac{E[RT]E[L]exp(\rho_{RT}\sigma_{RT}\sigma_L t)}{E[L]} \quad (2.35)$$

which can be simplified to

$$k = E[DA]exp(\rho_{DA}\sigma_{DA}\sigma_L t) - E[RT]exp(\rho_{RT}\sigma_{RT}\sigma_L t) \quad (2.36)$$

Note that $E[DA] \approx E[RT]$ since if one market was consistently higher or lower than the other, there would be arbitrage opportunities. The RT volatility σ_{RT} is usually higher than the DA volatility σ_{DA} due to real-time mismatches in supply and demand. I believe ρ_{DA} would be higher than ρ_{RT} since the load is usually cleared on the DA market and would have a stronger correlation with DA prices. In practice, k would be very close to 0, which is why I believe the team does not charge anything for energy pass-through deals. Nevertheless, this provides a theoretical framework to price the DART spread.

3. GREEKS, HEDGING AND MARGIN

3.1. Shape Swaps. In all of our analysis thus far, we have assumed that load and price are both random variables. However, a small portion of the FR portfolio consists of shape swaps, where the load is deterministic. For simplicity, let us consider a "1-hour" "shape" swap, where the load is known to be L . The payoff f from this swap would be $(k - P)L$. Solving for k by setting the expected payoff to 0 gives us $E[P]$ which is essentially just the forward. The same result can also be obtained by substituting a zero covariance in (2.1).

We solve for k and set that to be the price of the swap at initiation, but the payoff f varies as $E[(k - P)L]$, which gives us $f = kL - FL$, since L is deterministic (note that $F = E[P]$). The delta (Δ) of the swap can be calculated as

$$\Delta = \frac{\partial f}{\partial F} = -L \quad (3.1)$$

As one would expect, the Δ is negative since we are shorting a swap, and this can be delta-hedged by taking a long futures position with a volume of L as calculated in (3.1).

Let us now extrapolate this to a realistic shape swap. Let us consider a month-long shape swap where the load for all n hours of the month is deterministic and known to be L_1, L_2, \dots, L_n . We want to calculate the strike k such that the expected payoff f equals 0.

$$f = E \left[\sum_{i=1}^n (k - P_i) L_i \right] = 0 \quad (3.2)$$

Solving for k by setting $f = 0$, we get

$$k = \frac{\sum_{i=1}^n L_i E[P_i]}{\sum_{i=1}^n L_i} \quad (3.3)$$

On the FR desk, we use hourly shapers to model the prices for a shape swap. Let s_1, s_2, \dots, s_n be the hourly shapers for the month in question, with a forward price of F . (3.3) can then be rewritten as

$$k = \frac{F \sum_{i=1}^n L_i s_i}{\sum_{i=1}^n L_i} \quad (3.4)$$

As before, once we solve for k and set that to be the strike at initiation, the payoff can be calculated as

$$f = E \left[\sum_{i=1}^n (k - P_i) L_i \right] = \sum_{i=1}^n k L_i - F \sum_{i=1}^n s_i L_i \quad (3.5)$$

The delta of the shape swap Δ can then be calculated from (3.5) as

$$\Delta = \frac{\partial f}{\partial F} = - \sum_{i=1}^n s_i L_i \quad (3.6)$$

From (3.6), we see that the delta-neutral hedge volume of a month-long shape swap is the sum of the loads for the month weighted by the shapers. As long as the shapers do not change, the delta of the shape swap can be hedged even if the forward curve F changes. As stated before, the delta of a short position in a shape swap is negative, which can be offset by a long position in futures.

3.2. Variable volume swaps (load-following deals). The bulk of the FR portfolio consists of load-following deals or variable-volume swaps for which we have already derived the risk-neutral valuation price k . By substituting (2.20) in (2.1), we get

$$k = E[P] \exp(\rho \sigma_P \sigma_L t) \quad (3.7)$$

As before, once we solve for k and set that to be the strike or contract price at initiation, the payoff can be calculated as

$$f = E[(k - P)L] = kE[L] - E[PL] = kE[L] - E[L]E[P] \exp(\rho \sigma_P \sigma_L t) = kE[L] - FE[L] \exp(\rho \sigma_P \sigma_L t) \quad (3.8)$$

The delta of the variable-volume swap Δ can then be calculated from (3.8) as

$$\Delta = \frac{\partial f}{\partial F} = -E[L] \exp(\rho\sigma_P\sigma_L t) \quad (3.9)$$

(3.9) gives us a very important result. To delta-hedge a variable volume swap, it would be incorrect to naively hedge at the expected volume. Since ρ tends to be positive on the FR desk, we would need to "over-hedge" by a factor $\exp(\rho\sigma_P\sigma_L t)$. Note that this factor is greater than 1 if $\rho > 0$ and lesser than 1 otherwise, since the other terms inside the exponential are always positive. Note that the FR desk seems to follow a fairly static hedging strategy while the formulation above indicates that the portfolio needs to be rebalanced frequently to hedge delta. Needless to say, transaction costs also need to be taken into account when determining the optimal hedging strategy.

Since the delta is independent of F , the gamma γ of the swap which is $\frac{\partial \Delta}{\partial F} = 0$. The vega of the swap ν can be calculated from (3.8) as

$$\nu = \frac{\partial f}{\partial \sigma_P} = -FE[L]\rho\sigma_L t \exp(\rho\sigma_P\sigma_L t) \quad (3.10)$$

Note that the variable-volume swap has negative vega ν . To hedge this, we would require an instrument with positive vega, such as an option. As mentioned in the introduction, since a variable-volume swap has optionality embedded in it, we would require a combination of swaps/futures and options to hedge the important greeks namely delta and vega. As an exercise, the reader is also encouraged to calculate the other greeks rho ρ and theta θ although these are typically not hedged.

3.3. Margin. The FR desk has chosen to calculate the margin of a deal based on a certain risk-adjusted return on capital (RAROC) of the risk to the deal. The risk has been modeled as the tail Earnings-at-risk (T-EaR) which is explained below.

The earnings per MWh λ of a 1-hour load-following deal would be

$$\lambda = \frac{(k - P)L + \Delta(F_T - F_0)}{E[L]} \quad (3.11)$$

where we assume that the deal is made delta-neutral only at initiation with a volume Δ of long positions in futures/swaps struck at a price F_0 which is the futures price at initiation. Clearly λ is a random variable, that can take both



FIGURE 1. Hourly, daily, and alternate volatility by month for PJM-W

positive and negative values. Let $g(\lambda)$ be the probability density function of λ . The T-EaR is defined as follows.

$$TEaR = \int_{-\infty}^{\lambda_p} \lambda g(\lambda) d\lambda \quad (3.12)$$

where λ_p is the 5th percentile (P-5) earnings (λ). In other words, the T-EaR is the average of all the earnings (λ) outcomes that are equal to or lower than the P-5 earnings. The mathematically adventurous reader is encouraged to attempt to arrive at a closed-form solution for the T-EaR. Once the T-EaR is calculated, the margin is calculated as follows:

$$Margin = \frac{-TEaR \times RAROC}{1 - Tax Rate} \quad (3.13)$$

4. EMPIRICAL RESEARCH

4.1. Calculation of volatility. As seen in 1, I calculate the volatility of PJM-W LMPs using the following three approaches.

- a. log return of prices hourly divided by the square root of 1 hour (converted to days)
- b. log return of prices when aggregated to a daily level, divided by the square root of 1 day (converted to years)
- c. Alternate approach (the roundabout way the FR desk calculates load volatility)

The reason a, b, and c are different is as follows:

- i. First, the hourly vol and daily vol will not match. The hourly vol will generally be higher than the daily vol.
- ii. In hindsight, the alternate approach I followed was incorrect. I calculated the standard deviation to average ratio of hourly data (also called the coefficient of variation) every month, and did the square root of the log of 1 + the coefficient of variation squared divided by the square root of 15 days (converted to years).



FIGURE 2. Step-wise GBM




FIGURE 3. T-0 GBM

Not sure why I tried this first, but this is neither correct for step-wise nor T-0 GBM (explained in the next section).

Another learning from this analysis is that the hourly volatility for electricity prices is very high (10 to 20, as seen on 1). The volatility the FR desk receives from the options desk is not hourly but daily instead.

4.2. Types of Geometric Brownian Motion. In 2 and 3, I simulate correlated loads and prices using the same simulation inputs (load volatility = 0.25, price volatility = 1.1, and load-price correlation = 0.8), but the paths followed are very different. For path-independent derivatives like European options or “1-hour” variable-volume swaps, both types of GBM lead to the same valuation. However, for path-dependent derivatives such as Asian or lookback options, step-wise GBM is the appropriate right simulation approach.

4.3. Testing real-world prices for GBM. This experiment is inspired by chapter 12 in [2]. According to GBM, the log returns of prices should be IID and normally distributed. As a start, I look at the histogram of the log returns of a year of PJM-W hourly prices in 4. The histogram looks bell-shaped but does not tell us anything about the dependence between the data. To test that, we classify the data into 4 states as follows:

- State 1: Log return ≤ -0.1

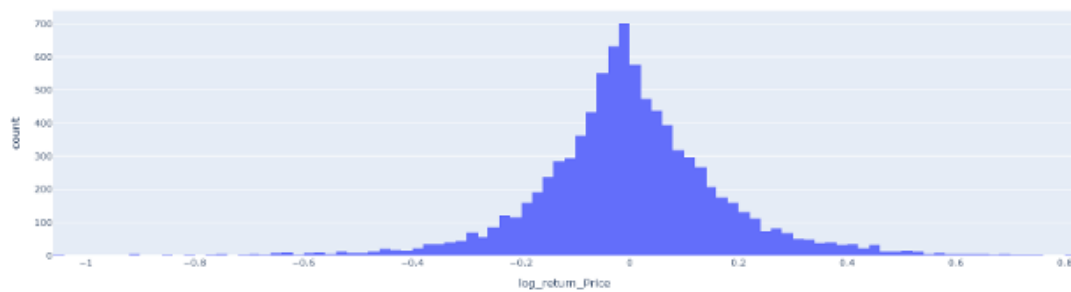


FIGURE 4. Histogram of 2025 PJM-W prices

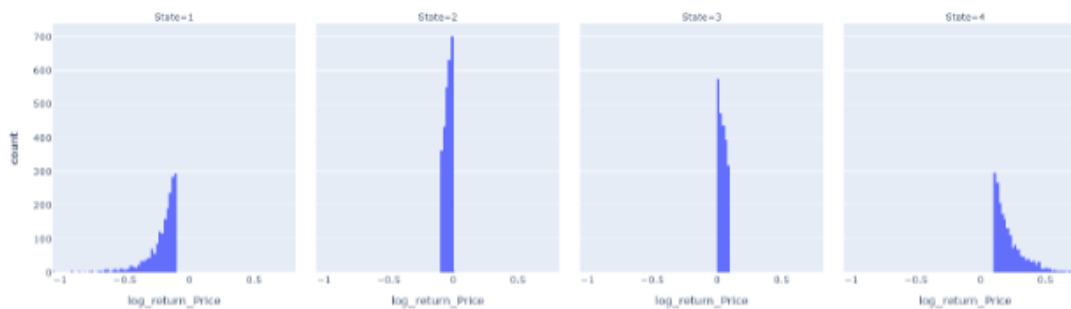



FIGURE 5. Histogram of prices belonging to each of the 4 states

FIGURE 6. Heatmap of the number of times state i is followed by state j

- State 2: $-0.1 < \text{Log return} \leq 0$
- State 3: $0 < \text{Log return} \leq 0.1$
- State 4: $\text{Log return} > 0.1$

The histogram of the data in the 4 states is seen in 5. To check for independence, we count the number of times state i is followed by state j . In 6 can see that the state 2 is frequently followed by state 2 (moderate drop followed by moderate drop) while state 4 is rarely followed by state 1 (large rise followed

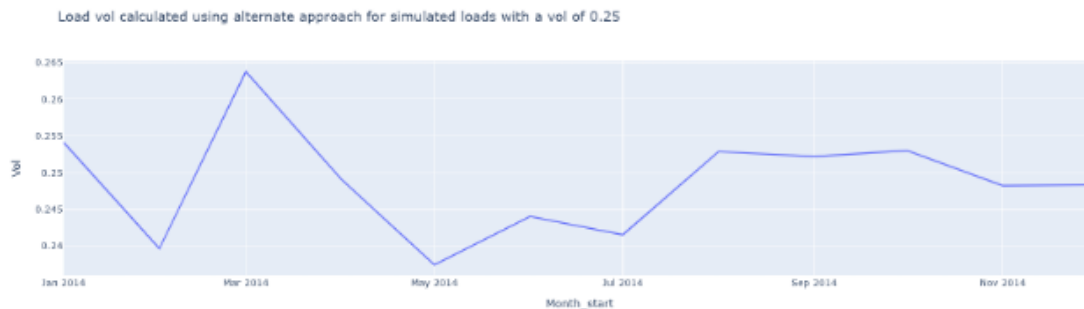


FIGURE 7. Alternative load volatility calculated for simulated loads

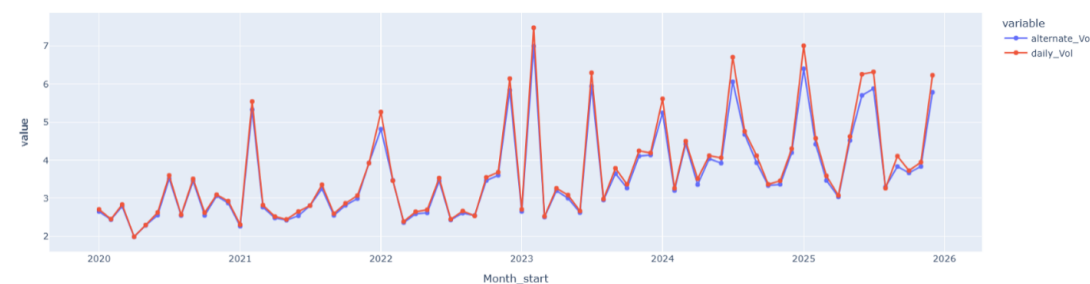


FIGURE 8. Comparison of volatility calculated using the definition and the alternate approach

by large drop). If the data indeed follows GBM, then the log returns should be independent and all states should roughly be followed by all states equally likely, but that is not the case here.

4.4. Verifying alternate volatility for simulated and real-world prices.

Using the T-0 GBM implemented in a previous section, I calculated the load volatility using the alternate approach where we use the standard deviation of load to calculate the load volatility. As seen in 7, the calculated load volatility closely matches the load volatility of 0.25 used to simulate the load in the first place. So, I can confirm that FR is using the right approach to calculate the load volatility. I also simulated GBM using the step-wise method and calculated the load volatility using the FR way (with some minor modifications), again leading to a similar result as shown in 7. Hence, I am confident of the method FR uses.

8 shows the same calculation as 7 with the only difference being that real-world PJM-W data was used for the analysis. 8 shows that the volatility calculated using both methods match each other closely.

As a final sanity check, I plotted the option-implied volatility by contract month and effective date in 9. It is interesting to note that the implied volatility of an option for a given contract month appears to increase as the option gets closer to maturity, likely due to higher trading activity.

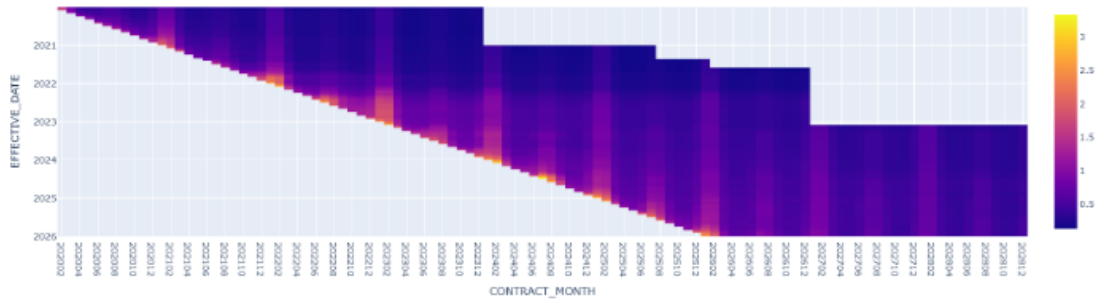


FIGURE 9. Option-implied daily volatility by contract month and effective date

Acknowledgement. I would like to extend a sincere token of gratitude to the Full Requirements team for their assistance in the review of this manuscript.

REFERENCES

1. Hull, John, 1946-. "Options, Futures, and Other Derivatives". Boston: Prentice Hall, 2012.
2. Ross, Sheldon M., 2011. "An Elementary Introduction to Mathematical Finance," Cambridge Books, Cambridge University Press, number 9780521192538, January.
3. Geometric Brownian Motion Implementation. `K:/Valuation/Analysts/Hemanth/Python Notebooks/Miscellaneous/Geometric_Brownian_Motion.ipynb`