INFORMS TutORials in Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Assets and Structured Hedges in Energy Markets: Severe Incompleteness and Methods for Dealing with It

Glen Swindle

TUTORIALS IN OPERATIONS RESEARCH
INFORMS 2016

 | ISBN 978-0-9843378-9-7
http://dx.doi.org/10.1287/educ.2016.0151

Assets and Structured Hedges in Energy Markets: Severe Incompleteness and Methods for Dealing with It

Glen Swindle
Scoville Risk Partners, New York, New York 10174, glenswindle@scovilleriskpartners.com

Abstract Risks in energy markets are inherently high dimensional because of large numbers of delivery locations and physical attributes, stochastic demand, and seasonality. By contrast, the number of instruments with sufficient liquidity to support hedging activities is relatively small, and it has never been able to span the set of risks sustained by market participants. This mismatch has spawned an interesting and arguably unique set of challenges related to the valuation and hedging of energy portfolios. Here, we will survey examples of such, including variable quantity swaps, generation, and structured asset hedges.

Keywords financial mathematics; energy markets; incomplete markets

1. Introduction

At first glance, energy markets resemble those of other asset classes. There are futures markets, which reflect pricing for delivery of energy over long tenors. Much as interest rate swaps markets can support the hedging of interest rate risk over tenors of 30 years or more, crude oil and natural gas futures markets serve a similar purpose for tenors in excess of 10 years. There are also spot markets and price indices to facilitate short-term purchase and sale of energy, analogous to cash markets in equities or daily LIBOR (London Interbank Offered Rate) fixings in rates. Similarities notwithstanding, energy markets have several important distinguishing features.

1.1. Historical Returns Volatility

Energy prices are typically much more volatile than prices in most other assets classes. Figure 1 shows the ratio of realized volatility of several reference assets spanning various asset classes versus that of West Texas Intermediate (WTI) crude oil futures (the "benchmark" reference crude oil price for North America with global relevance) on an annual basis. "SPY" is the S&P 500 exchange traded fund (ETF) price representing equities, the euro ("EUR") is the sample currency, the generic 10-year treasury ("GT10") was used for bond markets,[1] and the front-month gold futures contract ("GC") represented precious metals. On one occasion, the front-month gold contract exhibited realized volatility in excess of that of WTI, but such occurrences are clearly the exception, with most other ratios well below 50%.

[1] Returns volatility for GT10 was proxied by the product of duration, assumed to be 8.5 years, and change in yield.

Figure 1. Ratio of realized volatilities by calendar year for various assets to WTI.

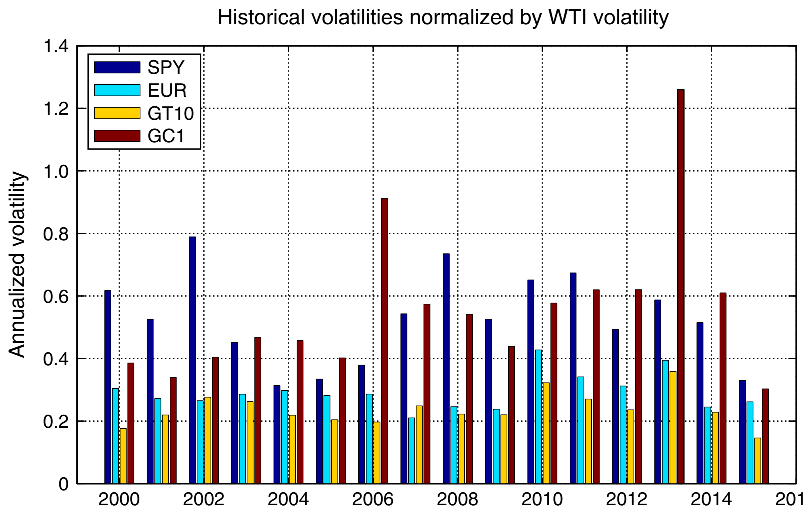

1.2. High-Dimensional Basis Risk

Another distinguishing feature of energy markets is the high-dimensional nature of basis risk. The term "basis" refers broadly to the price difference between any particular asset or commodity and that of an acknowledged reference asset—the "benchmark." All asset classes have such benchmarks, and portfolio management is often an exercise in managing a large set of technically distinct price risks with only a few benchmarks.

As an example, the typical Treasury bond portfolio will have positions in dozens of distinct bonds at any particular moment, but traders will usually use a small subset of liquid bonds or bond futures to rebalance their positions on a daily basis. The price spreads between the individual bonds and those of the liquid benchmark, often a bond futures contract, is an example of basis risk. In general, hedging with benchmarks works well since basis is typically stable, at least in comparison to energy.

In energy markets, basis trading refers to trading locational spreads—the difference between prices at distinct delivery locations. The number of such locations can be very large—U.S. natural gas has many dozens; U.S. power markets, thousands. Moreover, the dynamics of the price differentials between distinct locations can exhibit remarkably high variance. Figure 2 shows the historical relationship between two natural gas daily prices—one in the northeastern United States (TETM3) and one at the "benchmark" location, Henry Hub in Louisiana. The ratio of the two price series can routinely change by several multiples, in some extreme cases by an order of magnitude. There is also visible nonstationarity in behavior in the most recent years.

The fundamental purpose of commodities markets is to provide a mechanism for risk transfer from those with natural long positions (producers) to those with a natural short position (consumers). Because of the nature of their core commercial activities, these "naturals" have price risk at a myriad of idiosyncratic price points. By contrast, the most commonly traded instruments available for hedging are quite limited. In crude oil markets, the global benchmarks are the well-known WTI and Brent futures contracts. In U.S. natural gas markets, the Henry Hub futures contract is the benchmark, supplemented by a few relatively liquid regional hubs. The situation is similar in electricity markets, with only a handful of delivery locations that can legitimately be considered liquidly traded.

Figure 2. Locational returns of natural gas spot prices: TETM3 vs. Henry Hub (HH).

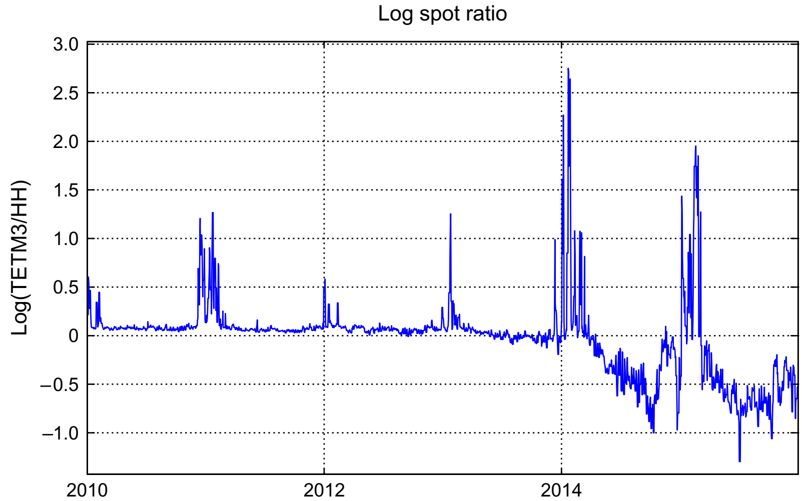

1.3. Modeling Implications

High-dimensional price risk and a low-dimensional set of functionally tradable instruments complicate what would already be challenging valuation problems. The situation has worsened in recent years as a result of the post-credit-crisis regulatory imperative, which has resulted in fewer liquidity providers (so-called dealers, which includes banks). Modeling approaches that explicitly address the lack of liquidity have become increasingly important. Here are some examples.

1.3.1. Project Finance. Energy infrastructure can be expensive to build. A typical large combined-cycle (efficient natural gas-fired) generation facility costs roughly one billion USD to build. Financing the development of such an asset usually involves significant borrowing by the builder or buyer—the sponsor. Lenders are well aware of the fact that, in the absence of a hedge, energy price volatility could cause the project to default on the loans. Therefore, it is very common for the lenders to require that the sponsor enter into hedges, often derivatives structures, designed to mitigate commodity price risk for the term of the loans.

Figure 3 is an illustration of how this works. The top set of cash flows represents the cost to build the asset and the cash flows that the asset will produce in the future—these

Figure 3. Project finance cash flows.

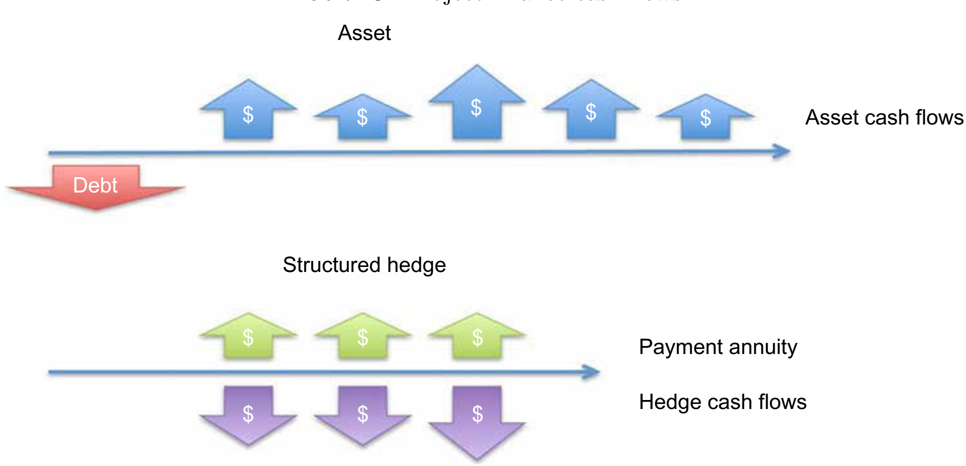

are random as a result of commodity price volatility as well as other risk factors. The lower set of cash flows represents an annuity that a hedge provider pays to the project in return for derivative payoffs that are intended to replicate the asset to a reasonable level of approximation. So long as the cash flows from the asset are closely matched by the payouts on the derivative, the annuity can be used to repay the loans, and everything works.

The picture is deceptive in its simplicity. There are three distinct problems embedded in this framework, each of which is challenging in its own right.

First, the asset cash flows must be modeled. Given a set of generation levels $Q(h)$ for a set of hours indexed by h, the asset yields a random set of discounted cash flows of the form

$$\Pi \equiv \sum_h d(h)Q(h)[p(h) - H_*(\cdots)g(h) - V], \tag{1}$$

where

—$p(h)$ denotes the spot price of electricity at the delivery node of the asset for hour h;
—$g(h)$ denotes the simultaneous natural gas price;
—$H_*()$ defines the asset conversion ratio ("heat rate"), which can depend on a variety of variables, including Q and environmental variables;
—V denotes additional costs, which can also depend on dispatch and environment; and
—$d(h)$ is the relevant discount factor.

Here, we have indexed all variables by an hour h, avoiding the notational complexities associated with the fact that gas prices are daily, and cash settlement, and hence discounting, is usually monthly.

The goal is to choose Q so as to maximize the risk-neutral expected value of the asset subject to $\mathcal{F}_h$-measurablility and consistency of Q with engineering constraints of the assets, which we will denote by $Q \in \mathcal{A}$. Choice of Q is the moral equivalent of basic option exercise, and it is often technically challenging to compute because of the complexity of the constraints $\mathcal{A}$ for most physical assets. In practice, the stochastic control problem is often handled heuristically as a result of the daunting computational requirements of exact solution. This topic has occupied the attention of many researchers over the years (see, for example, Carmona and Ludkovski [10]).

The second problem, the analysis of the hedge, is usually similar to the first: valuing a variation of the cash flows in (1) representing the hedge. We will denote these cash flows by $\tilde{\Pi}$. The optimization of $\tilde{\Pi}$ is typically simpler in structure than for Π, since hedge providers are often constrained in the complexity of what they can transact, because of both technological limitations (booking the trade in existing risk systems) and higher capital costs for less transparent structures.

A more substantive distinction is that $\tilde{\Pi}$ often has underlying price processes $\tilde{p}$ and $\tilde{g}$ that are different from the prices processes underlying Π. Hedge providers usually prefer to transact at liquid price points—not at the idiosyncratic delivery node of any particular asset.

This spawns the third, and often overlooked, problem. Given that Π and $\tilde{\Pi}$ are distinct, how should one calculate the distribution of the residual risks $\Pi - \tilde{\Pi}$? The most common methods to evaluate Π and $\tilde{\Pi}$ do not lend themselves easily to the construction of joint distributions.

1.3.2. Variable Quantity Risk. In the generation example just discussed, the quantity of electricity produced each hour is not known in advance. The quantity process $Q(\cdot)$ depends on the evolution of the price processes and is the result of optimization of (1). In many situations, however, the quantity of commodity produced or consumed is random for purely econometric reasons—not merely a function of option exercise. The following are some examples:

1. Oil and natural gas producers typically have multiple producing fields with varied geographic locations, each of which has many wells and each with varied rates of production as a result of unpredictable decline rates over the life of the well.

2. Refineries and generators have variation in production above and beyond the optimization $Q(\cdot)$ discussed above, as a result of mechanical failure and variation in ambient temperature.

3. A wind farm consists of a collection of turbines at distinct price points. Each will produce random quantities as a result of variations in wind. Similar problems apply to solar and other forms of distributed generation and battery storage.

4. Consumption by end users of all energy products changes on both long timescales as a result of macroeconomic factors and on short timescales as a result of weather. People need more energy in periods of extreme temperatures, and they all need it simultaneously.

In each of these examples, the result is a random volume of production or consumption at multiple locations. Diversification mitigates some of these risks—distinct oil fields can be viewed as independent. However, common factors such as temperature or wind tend to correlate behavior across multiple sources and sinks, resulting in undiversified risk.

The basic structure for each of these is disarmingly simple in form:

$$\Pi \equiv \sum_k \sum_n d(n) Q_k(n)[p_k(n) - p_k^*(n)], \tag{2}$$

where

- L indexes delivery location, and n is a time index;
- $Q_k(n)$ is the (random) quantity of production, where $Q_k(n) < 0$ implies consumption;
- $p_k(n)$ is the spot price of the commodity produced at L at time n; and
- $p_k^*(n)$ is the cost of production; for consumption, it is the contracted price received from consumers.

The challenges in dealing with portfolios of this form are manifold. First, the typical market participant has positions in many locations L, with an ability to hedge efficiently at only a few benchmark locations. Second, the correlation between $Q_k(n)$ and $p_k(n)$ is usually negative. Greater production (consumption) is associated with a lower (higher) spot price $p_k(n)$, a convexity that results in lower value and only partially hedgeable risks.

Focusing for the moment on a single delivery location, the quantity $Q(n)$ is a random variable; viewed collectively, $\{Q_\cdot\}$ is a stochastic process. Econometric analysis yields estimates for $\bar{Q}(n) \equiv E[Q(n)]$, in addition to other statistical attributes. The typical (though not optimal) hedging practice is to sell the expected quantity in the forward markets, the terminal payoff of which is

$$\tilde{\Pi} \equiv \bar{Q}(n)[p(n) - F(n)], \tag{3}$$

where $F(n)$ is the forward price at which the hedge was transacted. The hedged portfolio payoff is the difference of (2) and (3):

$$\Pi - \tilde{\Pi} = [Q(n) - \bar{Q}(n)][p(n) - p^*(n)] + \bar{Q}_n[F(n) - p^*(n)]. \tag{4}$$

It is the first term that is interesting, the second term being constant. Demand is positively correlated with price so that the expected value of the first term is negative. Correlation is against the holder of this position.

If forwards and options on $Q(n)$ were actively traded, then (4) could be treated as a quanto, and some practitioners do in fact attempt to construct forwards and volatilities for $Q(n)$, as well as correlations with underlying prices. The fact is, though, that derivatives on Q are simply not traded—volumetric risk is uncommoditized, and other methods are required.

1.3.3. Modeling Goals. The purpose of a modeling framework is to make reliable inferences about the value and risk of an illiquid (exotic or bespoke) payoff from the market prices of traded instruments. The more the target payoff differs from the traded instruments used to calibrate a model, the greater the modeling risk and, therefore, the more important the choice of the model.

A great deal of modeling research in commodities has focused on handling structures that depart from traded markets purely in the mechanics of the payoff. For example, any bespoke feature applied to a set of standard vanilla options, such as caps on total payoffs or knockout features, constitutes structural modifications to payoffs that commonly trade. The generation structure (1) also falls into this category if the underlying prices are liquid.

Another form of inference arises when the target structure depends on underlying prices and indices that are not traded. The variable quantity structures discussed earlier are an example. If, in addition, the price index is at an illiquid delivery point, then both underlying indices require inference of some sort. In situations such as this, a modeling framework should do more than simply provide point estimates of values and optimal hedges. It should also yield distributional results characterizing the unhedgeable residuals risks.

2. Basic Concepts and Methodologies

Before plunging directly into methods of analyzing energy portfolios, some basic features of energy markets warrant review.

2.1. Forward Markets and Spot Prices

Commodities markets "trade forward." By this, we mean that the most commonly traded instruments are for delivery of the commodity at a defined time in the future, at a defined delivery location, and of a precise specification in return for a fixed price.

Figure 4 shows the forward curve for WTI crude oil delivered in Cushing, OK. Each point on the curve represents the settlement (end-of-day) price on January 15, 2016 in dollars per barrel for delivery in the corresponding month.[2] The WTI contract involves physical delivery. If you have purchased such a contract and hold it to the delivery month, you will receive barrels of oil.

In what follows, we will denote the forward price of a commodity at a current time t for delivery at a future time T by $F(t,T)$.

At delivery, the commodity is sold on a daily basis at what is termed the "spot" price, meaning the price of a commodity for delivery "now." Formally, the spot price is $F(t,t)$ in our notation. In contrast to forward prices, the meaning of which is unambiguous, the concept of a spot price is actually somewhat slippery. In crude oil and natural gas markets, spot prices are provided by vendors (for example, Platts and Argus) who survey market participants for transactions involving (nearly) immediate delivery during a trading day. From such data the vendors provide their estimates for where the commodity traded during the day. The data behind Figure 2 used exactly this price data type. The integrity of these surveys has always been a subject of discussion, if not at times outright controversy. By contrast, many electricity markets are administered by nonprofit organizations that use well-defined optimization algorithms to match generation and load on an hourly basis, in the process producing shadow prices that are the spot prices of the hundreds or thousands of delivery nodes in the system.

Two important points warrant mentioning. First, the forward curve defines the prices at which you can transact now for future delivery periods. It is not, however, a particularly useful predictor of future spot prices. Figure 5 shows a sequence of Brent (the global

[2] Technically, these are futures prices for WTI contracts trading on the CME/NYMEX exchange. We will not delve into the technical and operational differences between futures and over-the-counter forward transactions here.

FIGURE 4. WTI forward curve.

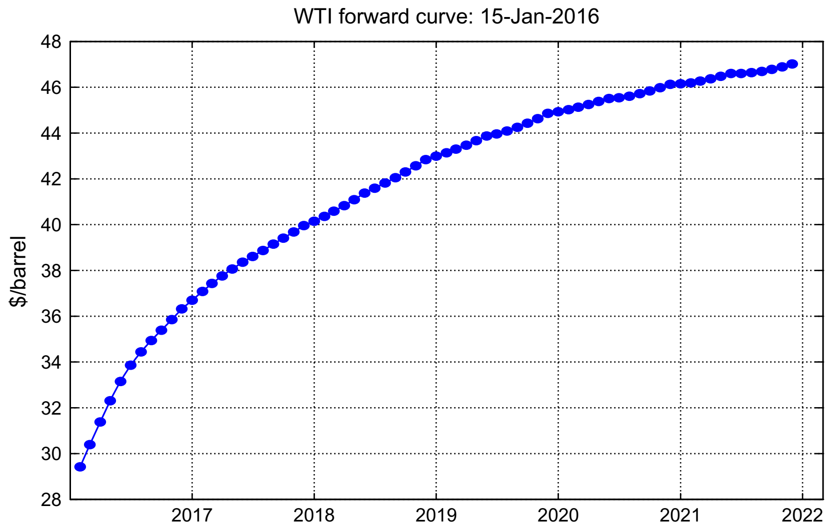

FIGURE 5. Brent forward curves with spot overlay.

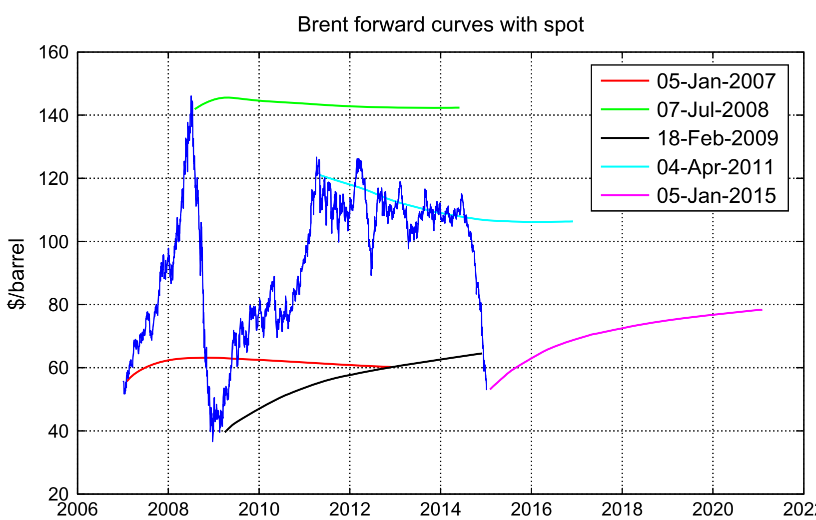

crude oil benchmark) forward curves with subsequent overlaid Brent spot prices. There is little relationship between the forward price at a particular time and the spot prices that eventually prevailed.

Second, the forward curve can be viewed as a yield curve by defining the forward yield as

$$y(t,T,T+S) = \frac{1}{S}\log\left[\frac{F(t,T+S)}{F(t,T)}\right]. \tag{5}$$

This is the yield that could be achieved by purchasing the commodity for delivery at time T and selling it later at time $T+S$—provided that the buyer has the facilities required to store the commodity in the interim. When forward curves are upward sloping (this is called "contango"), forward yields are positive. If the forward yield exceeds borrowing costs, the effect is to incentivize those with storage facilities to add to inventory, thereby relieving excess supply. Conversely, when yields are negative (referred to as "backwardation"), those with inventory are incentivized to sell it, thereby relieving a supply shortage. Figure 6

FIGURE 6. WTI forward yields vs. OECD inventory.

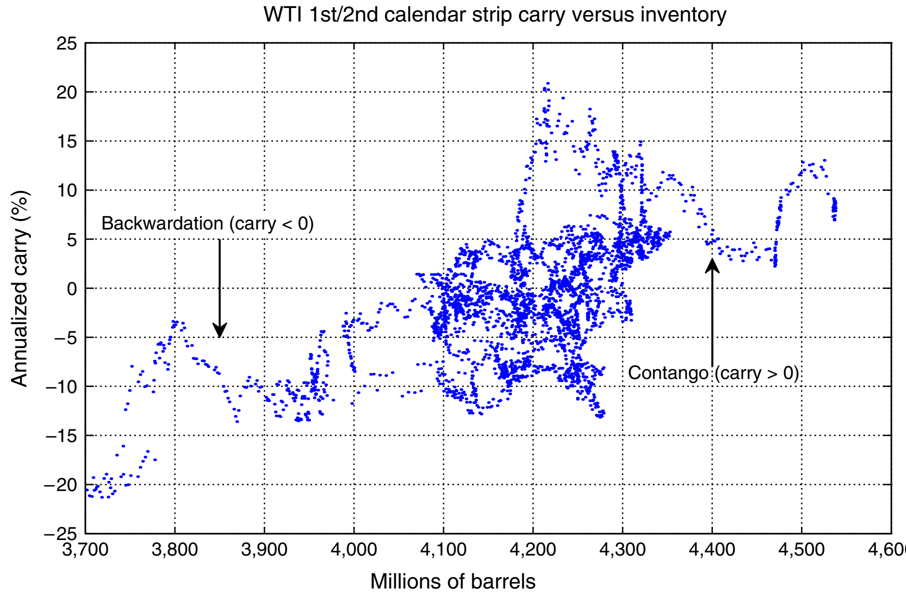

illustrates this relationship; high Organisation for Economic Cooperation and Development (OECD) crude oil stocks are associated with high WTI forward yields, and vice versa. This phenomenon is observed broadly across consumable commodities markets.

2.2. Options Markets

Options prices, when available, provide the implied volatilities used to calibrate models. Three types of options commonly trade; these are differentiated by the timescale of delivery.

Monthly Options: The most commonly traded energy options, monthly options exercise into delivery for a calendar month. The option payoff is a function of $F(\tau, T_m) - K$, where T_m indexes monthly delivery, $F(\tau, T_m)$ is the value of the futures contract at expiration/exercise τ, and K is the strike. Exchange traded options of this type often have American exercise features.

Swaptions: These are options that exercise into a strip of contracts at a fixed strike K. Often the strip is one calendar year. At expiration, the value per unit notional is

$$V_\tau = \frac{1}{12} \sum_m d(\tau, T_m)[F(\tau, T_m) - K], \tag{6}$$

where τ is the exercise time.

Daily Options: Options involving daily exercise are traded in power and natural gas markets. A daily option trade is, in fact, a set of distinct options for each delivery date in a month. For example, a call option would settle on the value $\sum_{d \in m} \max[F(T_d, T_d) - K, 0]$.

The options that trade in a given energy market are usually an order of magnitude larger in timescale than the inherent risks in the markets. Oil and natural gas have daily price risk, but by far, the most liquid options are monthly in nature. For power, where spot prices are printed hourly (or more frequently in some markets), the lowest timescale for commonly traded options is daily. This gap in timescales is a fundamental source of model risk.

FIGURE 7. WTI realized returns volatilities by tenor (2009–2015).

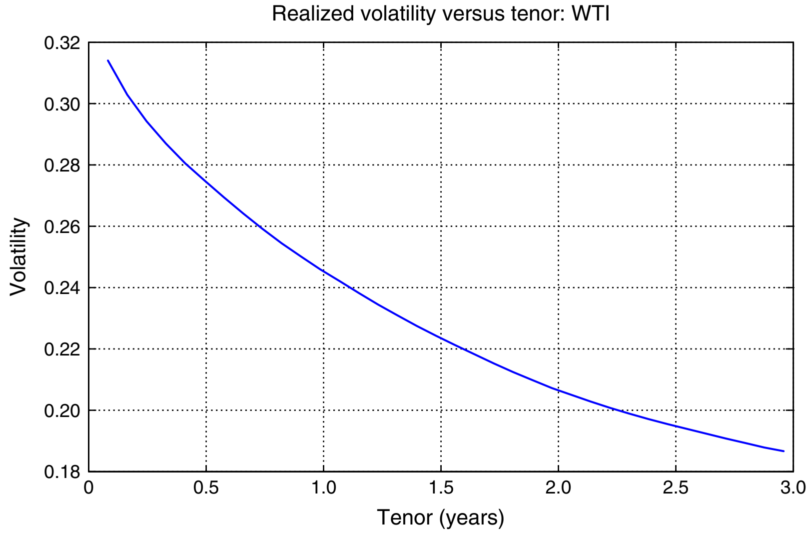

2.3. Price Dynamics

Each energy market has its own idiosyncrasies with respect to price dynamics as well as what types of structures are traded. Model selection involves the usual dialectic—the trade-off between ease of application and stability versus the complexity required to accommodate observed features of price dynamics. We survey a few facts about energy price dynamics that are generally viewed as important.

2.3.1. Volatility Decay. Returns volatility of forward prices decays systematically with the tenor of the contract. There are exceptions to this statement, especially for commodities that exhibit seasonal structure such as some refined products, natural gas, and electricity. However, seasonality aside, this feature is common across consumable commodities. To illustrate, Figure 7 shows realized returns by tenor for WTI futures between 2009 and 2015 annualized. A contract for delivery in two years sustained a volatility that was roughly two-thirds of that for a contract with delivery in one month.

Traders are, of course, aware of this behavior. Figure 8 shows the at-the-money (ATM) implied volatilities for WTI monthly options.[3]

2.3.2. Skew. Empirical price returns are, in general, not normally distributed. Returns statistics for rolling (constant maturity) price returns for crude oil futures are nonnormal at all tenors. The same pertains to other energy asset classes as well, with even more dramatic departures from normality at daily or hourly timescales.

Defining returns for spot prices is not a trivial matter. For commodities that can be difficult to store (electricity being the canonical example), the naïve approach of calculating returns on the daily spot price p: $\log[p(d+1)/p(d)]$ yields results that are of dubious relevance: a commodity delivered on day d cannot be carried to day $d+1$. Different delivery days are in effect two different assets, and such a return series is meaningless for trading and hedging.

A metric that is more sensible references the daily spot price $p(d)$ to the futures price for the delivery month at the time the contract expired just before the delivery month:

$$r(d) \equiv \log\left[\frac{p(d)}{F(T_e, T_m)}\right], \tag{7}$$

[3] The apparent anomalies at longer tenors are “noise” in the settlement prices and would typically be smoothed by in-house systems before use.

FIGURE 8. WTI ATM implied volatilities.

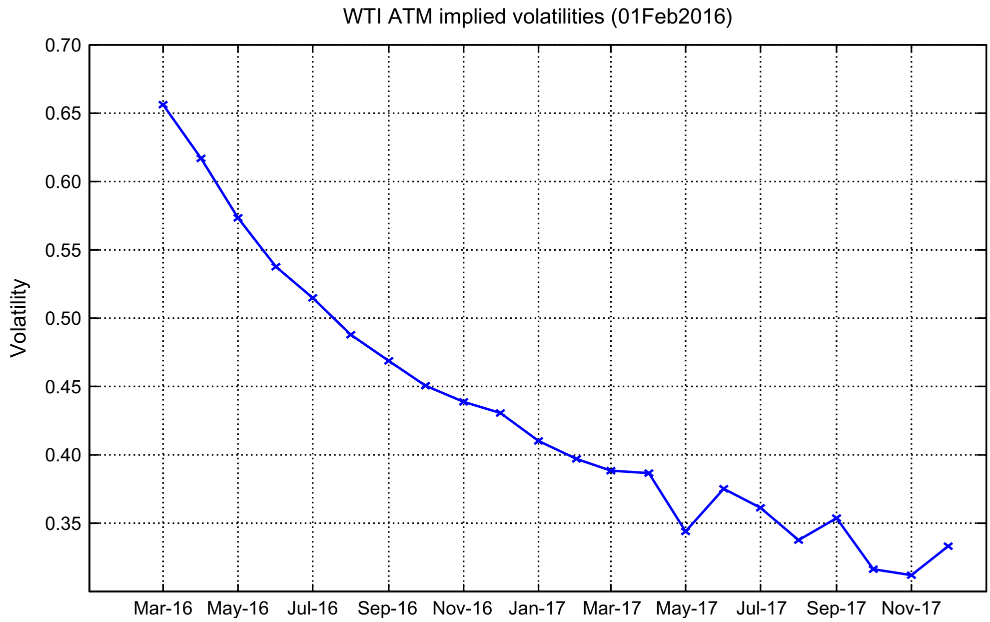

Source. Bloomberg.

FIGURE 9. Henry Hub natural gas spot returns statistics.

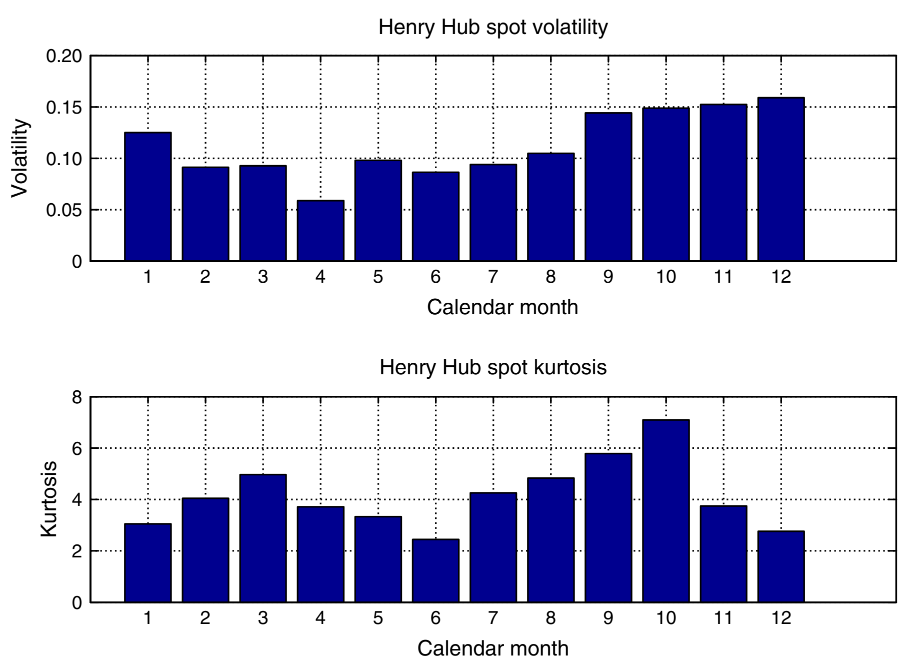

where $F(T_e, T_m)$ is the expiration price of the futures contract for the month in which $d \in m$. Figure 9 shows the average volatility and kurtosis by calendar month for the natural gas benchmark location Henry Hub calculated from 2004 through 2015. The seasonal structure is due in part to the confluence of winter demand and hurricane season, which peaks in early autumn in North America. The following items are of particular note:

—The volatilities are not annualized. Annualized, these would be roughly five times greater, showing the relatively high levels of returns volatility at short timescales.

—Kurtosis can be quite high in some months, well exceeding the normal kurtosis of 3.

FIGURE 10. PJM Mid-Atlantic load vs. KPHL daily average temperature.

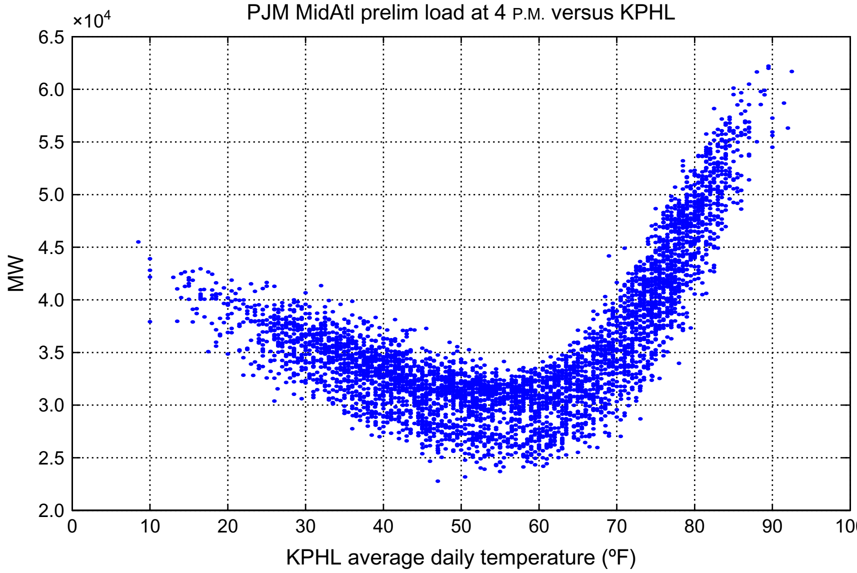

2.3.3. Weather Dependence. Energy has two dominant uses. The first is to move things—transportation. The second is to control the ambient temperature that we experience through heating in the winter and air conditioning in the summer. The impact of weather as a demand and price driver is visible in many ways. As one example, Figure 10 shows electricity demand in a defined subset of the PJM electricity market in the eastern United States at 4 P.M. versus the average daily temperature at a weather station in Philadelphia (KPHL). Among practitioners, analysis of variable demand structures often requires explicit modeling of temperature as an underlying stochastic driver.

3. Modeling Frameworks

Modeling in energy is an exercise in capturing the features of price dynamics relevant to any particular problem while maintaining tractability. The main issues are as follows:

—Can the models be calibrated to forwards and traded options?

—Are price dynamics of the model consistent with the statistical features of realized prices?

—Are the models of a form that can readily be coupled to underlying drivers such as temperature and related demand indices?

—Can the framework be used when the underlying indices are illiquid or simply untraded?

—Are the models tractable, allowing efficient valuation and calculation of Greeks?

3.1. Martingales and Black-76

Forward and futures prices are martingales—forward prices under the *T-forward* measure, futures under the cash measure.[4] The martingale property is due to the absence of a cash cost of carry and is a modeling constraint that is handled one of two ways. Forward prices $F(t,T)$ can be explicitly represented as martingales. Alternatively, a spot price process p_t

[4] The mechanics of forwards and futures are distinct. However, aside from credit risk, the two are functionally the same from a valuation perspective. See, for example, Eydeland and Wolyniec [16], Geman [17], or Swindle [29] for further discussion of the martingale constraint on futures models.

can be constructed from which $F(t,T) = \tilde{E}[p_T \mid \mathcal{F}_t]$ is a martingale; here, $\tilde{E}$ is the (or a) risk-neutral measure and $\mathcal{F}_t$ is the usual filtration.

The simplest approach to risk-neutral valuation in commodities is that of Black-76 (Black [7]), in which a drift-free geometric Brownian motion generates returns:

$$\frac{dF(t,T)}{F(t,T)} = \sigma\, dB_t. \tag{8}$$

Integration of (8) yields

$$F(t,T) = F(0,T)e^{-(1/2)\sigma^2 t + \sigma B_t}. \tag{9}$$

The implication is that the spot price process is governed by

$$p_t \equiv F(t,t) = F(0,t)e^{-(1/2)\sigma^2 t + \sigma B_t}, \tag{10}$$

or, equivalently,

$$dp_t = \left.\frac{\partial F(t,T)}{\partial T}\right|_T dt + \sigma\, dB_t. \tag{11}$$

The spot price drifts at the local slope of the prevailing forward curve.

Black-76 is the first of a class of models commonly referred to as *reduced-form* models. The volatility σ can be calibrated to the price of a single option on $F(t,T)$. However, while dealing with carry properly, none of the stylized features is captured, and there is little to recommend this as a useful model for anything other than the simplest valuation problems.

3.2. Reduced-Form Models

Departing from chronological development of more refined models, Gaussian multifactor models constitute a nontrivial extension of Black-76 and one that is in common use. These take the form

$$\frac{dF(t,T)}{F(t,T)} = \sum_{j=1}^{J} [\sigma_j(t,T)e^{-\beta_j(T-t)}\, dB_s^{(j)}]. \tag{12}$$

The underlying Brownian motions can, in general, be correlated. The approach, first articulated in this form by Clewlow and Strickland [13], is clearly analogous to Heath–Jarrow–Morton (HJM)-style rate models but of a simpler form as a result of the absence of drift.

This class of models offers a considerable degree of flexibility. The set of exponential basis functions for local volatility can closely approximate empirical volatility surfaces (Swindle [29]), it is by construction consistent with forward markets, and it can be calibrated to match implied volatilities of options of different delivery periods in the multifactor setting. Another advantage is that the exponential form of the factors yields a low-dimensional state representation as a result of the common appearance of the integrals of the form $Y_t = \int_0^t e^{-\beta(t-s)}\, dB_s$ in (12). These processes are Ornstein–Uhlenbeck diffusions $dY_t = -\beta Y_t\, dt + \sigma\, dB_t$, which can be integrated in closed-form and permit finite-dimensional representation of stochastic control problems.

In the early stages of interest rate modeling, the focus was almost exclusively on short-rate dynamics. The same pertained to early commodities price models, with specification of spot price processes being the usual starting point. In fact, (11) hints at the near equivalence of the two approaches (see Baxter and Rennie [5], which elaborates on this point in the rates setting). One example is the Gibson and Schwartz form (Gibson and Schwartz [20]) of coupled diffusions:

$$\begin{aligned} dp_t &= (r_t - \delta_t)p_t\, dt + \sigma p_t\, dB_t^{(1)}, \\ d\delta_t &= \kappa(\theta - \delta_t)\, dt + \gamma\, dB_t^{(2)}, \end{aligned} \tag{13}$$

with $d\langle B^{(1)}, B^{(2)}\rangle_t = \rho$. The spot price p_t is a geometric Brownian motion with a drift term corresponding to financial carry at the short rate r_t modified by a "convenience yield" δ_t, itself a diffusion of Ornstein–Uhlenbeck form.

It turns out that the modeling class (12) subsumes (13)—a fact that is not obvious but can be established by direct calculation of $F(t,T) = \tilde{E}[p_T \mid \mathcal{F}_t]$. See Carmona and Ludkovski [9] and Swindle [29].

While (12) is much more flexible than Black-76, it does have significant weaknesses, the most significant of which is that returns are normal; skew cannot be handled explicitly. This limitation and the observed "spiky" behavior of spot prices (recall Figure 2) has spawned a large number of variations of reduced-form models. These fall into distinct categories:

- random or stochastic volatility models—see, for example, Hikspoors and Jaimungal [22], Ribeiro and Hodges [26], and Trolle and Schwartz [30];
- jump diffusion models such as those discussed in Barz and Johnson [4], Cartea and Figueroa [12], Deng [15], Geman and Roncoroni [19], Kaminski [23], and Kholodnyi [24]; and
- Lévy processes—see Barndorff-Nielsen et al. [3], Benth et al. [6], and Geman [18], and references therein.

3.3. Application of Reduced-Form Models

Deciding what class of reduced-form models to use is an exercise in balancing tractability and consistency empirical features of price dynamics. For relatively liquid price points, the procedure is usually as follows:

- The market data to which the model will be calibrated are identified. This always includes the forward prices and will also typically include some or all of the variety of options that are traded.
- The set of parameters that will be varied in the calibration process is selected. For example, for the Gaussian multifactor models (12), the local volatilities σ_j are usually adjusted to match option prices.
- The remaining parameters are estimated statistically.

It is usually relatively simple variations of these models that are used in practice because of limitations in the market data to which model parameters can be calibrated and to the imperative that model calibrations be stable.

Parameter estimation is almost always required given the paucity of market data. This may seem similar to the econometric models discussed momentarily, but the key distinction is that additional stochastic drivers are not introduced at this stage. The parameter values used—for example, the mean-reversion rates in (12)—are point estimates.

The situation is similar when reduced-form models are applied at illiquid delivery locations. If a structure is being valued that settles on spot price p_t at an illiquid location, the usual approach is to estimate the spread between p_t and a spot price $\tilde{p}_t$ with greater liquidity. The estimate is then treated as a constant in the subsequent valuation, with basis risk between p_t and $\tilde{p}_t$ not explicitly considered.

Reduced-form models are most useful when underlying prices are liquidly traded. Econometric and structural models, to which we turn next, are designed primarily to handle situations in which liquidity is wanting.

3.4. Econometric and Structural Models

In econometric and structural modeling, the dominant consideration is consistency with empirical price dynamics. Econometric models posit functional forms for spot price dynamics in the physical measure, the parameters of which are estimated via regression. The form of the regression used is often based on "known" attributes of the market. Structural models are best described as turbo-charged econometric models, adding far more detail about market structure to the design of parametric models.

FIGURE 11. Barlow stack and load distributions.

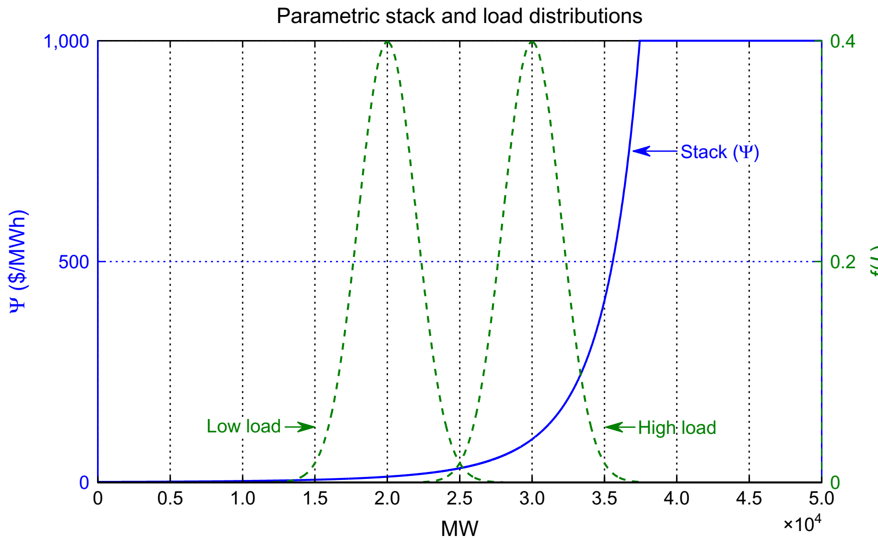

What do we mean by *known* attributes? In electricity markets, where these approaches are most developed,[5] spot prices are known to be set by well-defined algorithms taht minimize costs subject to reliability constraints. One view of price clearing is obtained by sorting all available generation in increasing order of cost—this is called the generation "stack." At any particular time, the spot price can be thought of as the value where demand intersects the price stack.

Figure 11 shows a version of this concept for one of the earliest electricity structural models introduced by Barlow [2].

In the absence of storage, market clearing is essentially an instantaneous affair with limited intertemporal coupling. The general structural model takes the form

$$p(t) \approx \Psi[L_t \mid \vec{F}_t], \tag{14}$$

where Ψ is a function that specifies the cost of supply given a set of input fuel prices $\vec{F}_t$. This is a mapping from demand L_t to spot price p_t.

Superimposed on the price stack in Figure 11 are two normal probability densities for future loads, corresponding to low and high cases. In each case the price distribution resulting from (14) will result in correlated demand and price processes and, depending on the form of Ψ, skewed distributions in high demand cases—both are observed and desirable features.

Many researchers have developed variations of parametric stacks models; see Aid et al. [1], Carmona and Coulon [8], Carmona et al. [11], Coulon et al. [14], and Skantze et al. [27, 28] for only a partial list, in addition to countless proprietary efforts by industry researchers. Structural models for storable commodities have also been a topic of interest to researchers for many decades. These are, however, used less in practice due in large part to daunting computational requirements. See Pirrong [25] for a thorough discussion of the potential strengths and limitations of such efforts.

As with more complex reduced-form models, calibration becomes increasingly difficult with model complexity. Econometric models take a step back from directly modeling caricatures of market clearing mechanics, using stylized facts to motivate the forms of regressions used. As an example, Figure 12 plots the ratio of daily PJM Western Hub[6] peak power prices to a locational natural gas spot price (TETM3) versus the natural gas spot price.

[5] Structural models have had the greatest impact in electricity markets, where the fact that electricity remains expensive to store (this may change, of course) simplifies the modeling requirements.

[6] Western Hub is the most liquidly traded location in the PJM market.

FIGURE 12. PJM Western Hub/TETM3 vs. TETM3.

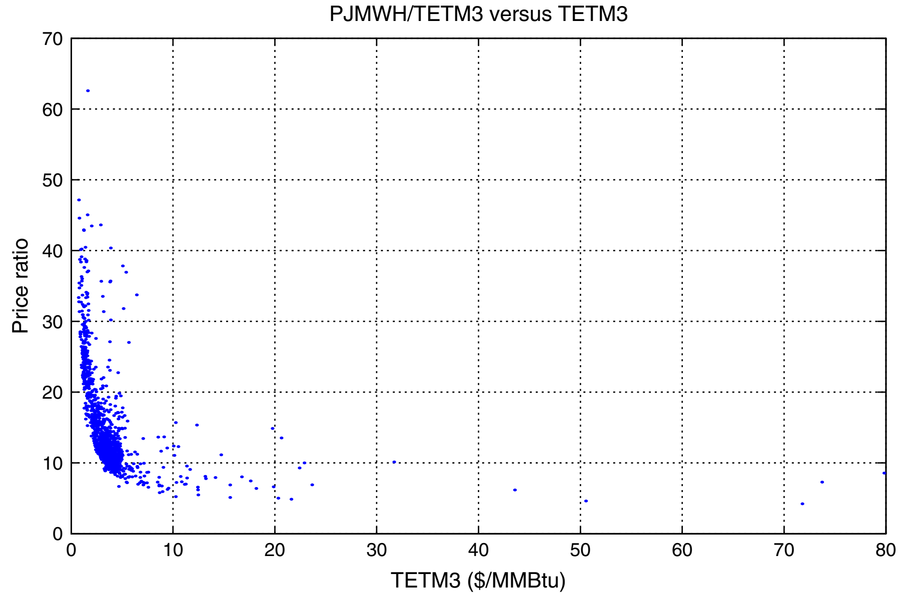

FIGURE 13. PJM Western Hub/TETM3 vs. KPHL.

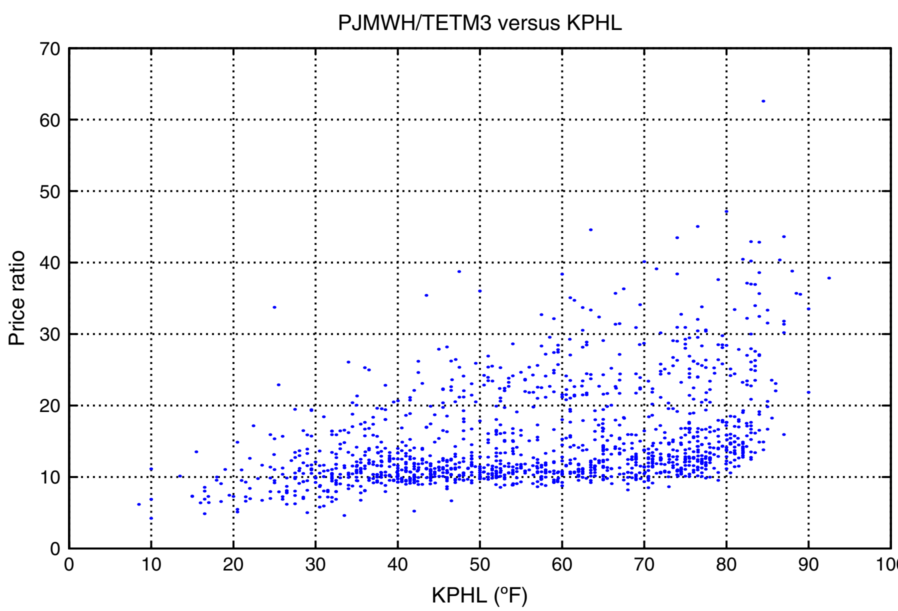

The expansion of this ratio as natural gas prices decrease is consistent with known behavior of electricity price stacks and market clearing. Figure 13 is another view of price dynamics, showing the same price ratio versus the daily average temperature at a related temperature index, KPHL (Philadelphia). The observable feature of higher price ratios at high temperatures is also consistent with the price stack (recall Figure 11).

These figures motivate regression forms that would be challenging to divine without some a priori knowledge of the system. In what follows we will use the form

$$\log\left[\frac{p(d)}{p_*(d)}\right] = \alpha + \gamma p_*(d) + \sum_{k=1}^{K} \lambda_k \theta(d)^k + \epsilon_d, \tag{15}$$

where $p(d)$ is the target spot price (PJMWH) and $p_*(d)$ is the reference price (TETM3). The function θ is a transformed temperature designed to avoid unrealistic simulated values when temperatures are outside of the historical data set:

$$\theta(d) \equiv \frac{e^{\lambda(d)}}{1+e^{\lambda(d)}}, \tag{16}$$

with $\lambda(t) \equiv (\tau(t) - \tau_{\text{ref}})/w$, and $\tau(d)$ is the daily average temperature (in this case, KPHL). The τ_{ref} and w are selected to be characteristic mean and width of temperatures realized over the historical data set.

The next step after regression estimation is to use (15) to construct spot prices under the physical measure via simulation. This means that we need to decide how to simulate the reference price $p_*(d)$, the residuals $\epsilon(d)$, and finally temperature τ_d.

Temperature plays a central role in this framework. Simulations are usually based on similar econometric models that exploit the existence of long quasi-stationary data sets to calibrate models of the form

$$\tau_d = \mu(d) + \sigma(d) X_d, \tag{17}$$

where X_d is a stationary process, and the mean is represented as

$$\mu(d) = \alpha_0 + \alpha_1(d - d_*) + \sum_{k=1}^{K} \big[c_k \sin(2\pi k\Phi(d)) + d_k \sin(2\pi k\Phi(d))\big]. \tag{18}$$

Here, d_* is a reference date and $\Phi(d)$ is the fraction of the year corresponding to d: $\Phi(d) = (d - \text{BOY}(d))/365$. This form accommodates both seasonality and drift, whether arising from urban growth or climate change (see Swindle [29] for more detailed discussion).

Figure 14 illustrates the results for KPHL. The selection of the number of Fourier modes to retain used an out-of-sample selection method.

We will discuss the remainder of the simulation process in the next section.

Figure 14. Calibration of temperature.

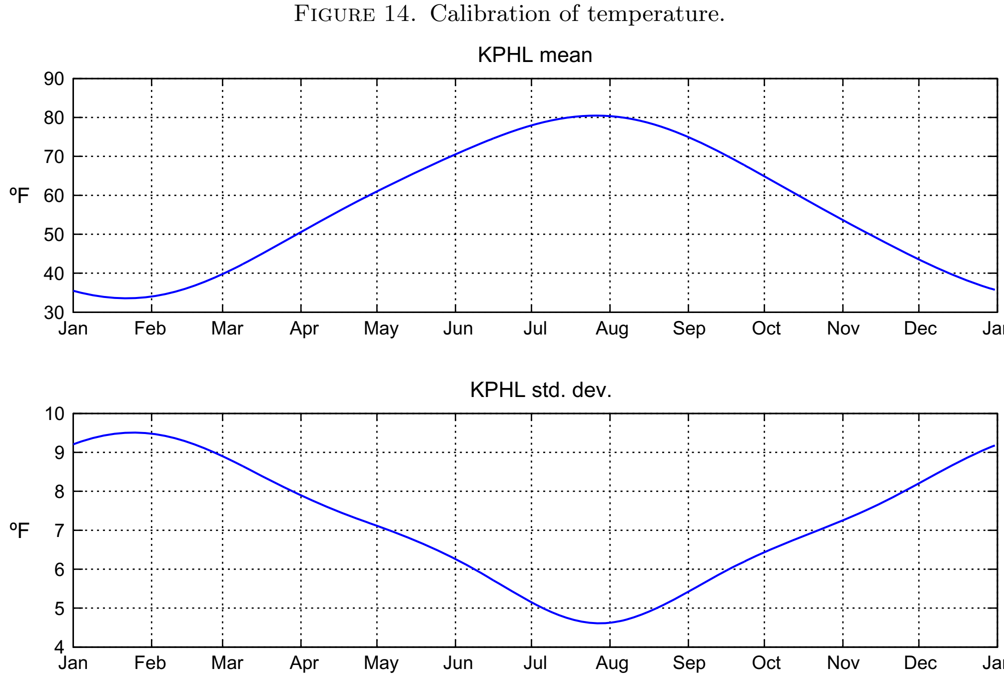

Figure 15. Hierarchical representation of multiple delivery locations.

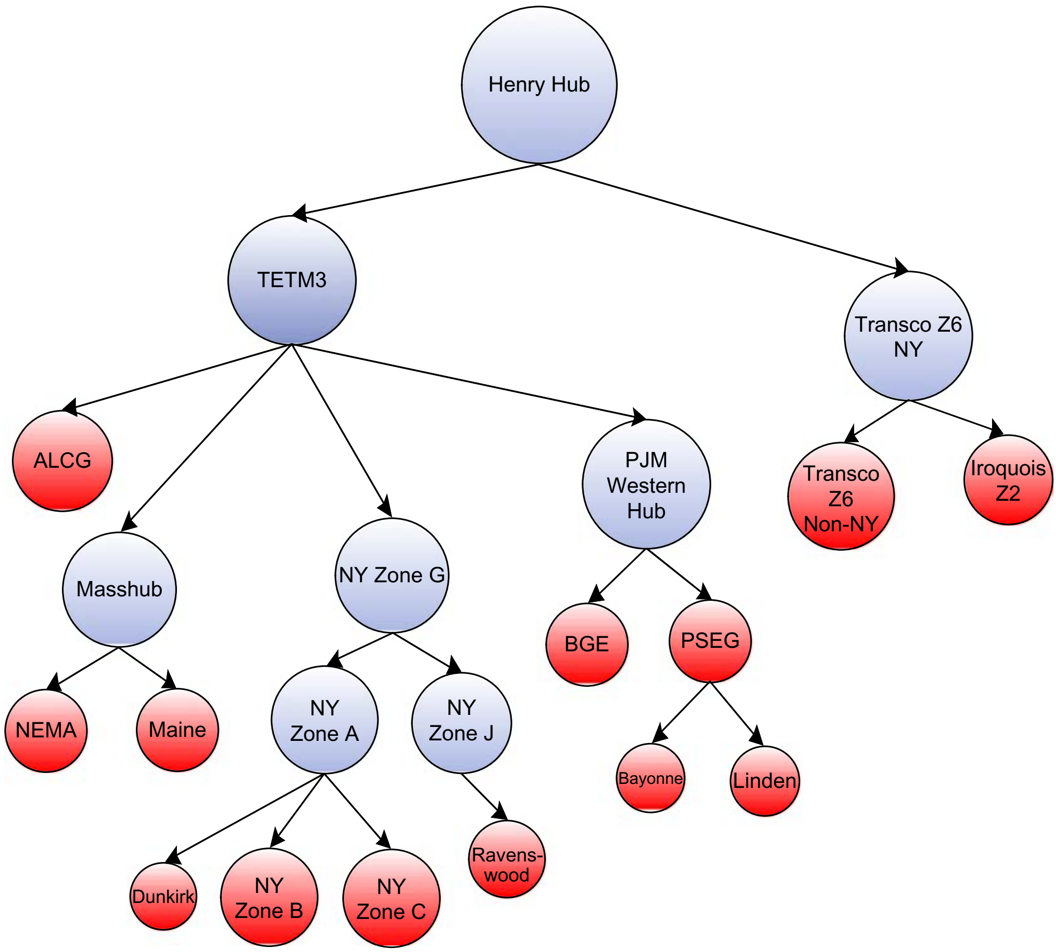

3.5. Application of Econometric Models

To calibrate and simulate econometric models across a broad swathe of delivery locations, some a priori organization is imperative. To this end, Figure 15 shows a sample hierarchy of prices starting at the natural gas benchmark location Henry Hub and descending hierarchically to delivery locations with lower levels of liquidity. This graph is intended to represent the sequence in which regressions of forms such as (15) are performed. Hierarchical organization of prices and risk is commonly used by traders in price marking and is also embedded in the design of some risk systems. From a stochastic modeling perspective, see, for example, Grine and Diko [21] and Swindle [29].

Once each model corresponding to each pair of nodes in the tree is calibrated, simulations are performed. To get things started, the benchmark natural gas contract must be simulated—subsequent nodes on the tree follow iteratively using the simulations of the parent node and the relevant temperatures. By definition, the benchmark is replete with liquid forwards and options markets. A common choice for the treatment of the benchmark is to use one of the reduced-form models discussed in the previous section. Econometric modeling frameworks are, therefore, effectively appendages to reduced-form models, expanding the set of tractable delivery locations.

In the results that follow, the benchmark Henry Hub prices are simulated using a two-factor version of (12) calibrated to NYMEX Henry Hub options. Then, given joint temperature

Figure 16. Residual correlation between weather stations.

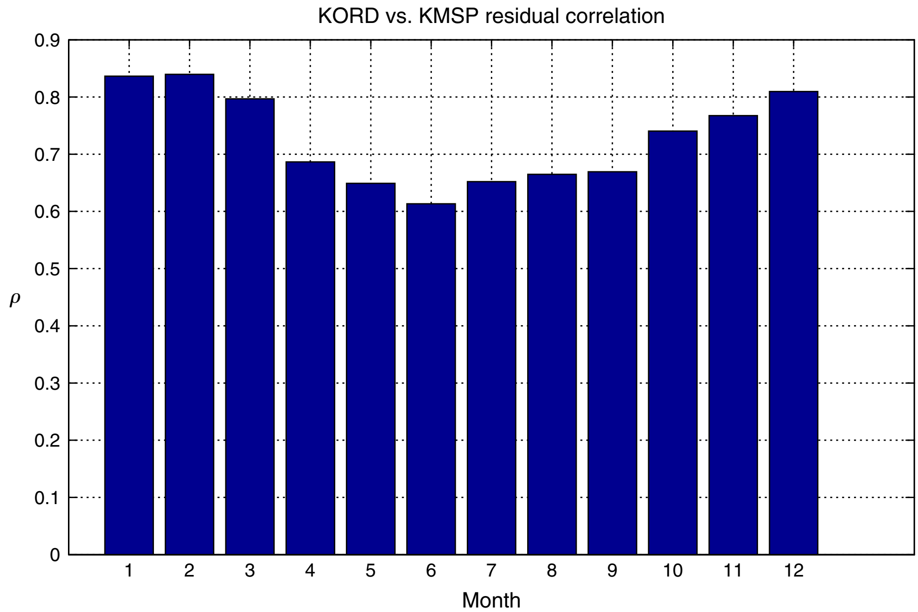

simulations across all relevant locations, spot prices are simulated at each node of the hierarchy by iterative applicate of (15). The result is a physical measure for spot prices at all nodes.

In the one-dimensional setting where a single temperature is driving a single spot price, a parametric representation (for example, ARMA models) of temperature residuals can be considered. However, attempting to simulate a large number of residuals series collectively is fraught with difficulties. Figure 16 illustrates these challenges. This plot shows the empirical correlation of residuals X_d in (17) between two relatively close weather locations, KORD in Chicago and KMSP in Minneapolis, by calendar month using roughly 40 years of data. There is pronounced seasonality. Moreover, this occurs to varying degrees between other weather locations, rendering a construction of a parametric representation a daunting task.

An alternative, and the one used in the results below, is to use a seasonal bootstrap to simulate temperatures and a bootstrap conditioned on simulated temperatures to resample the spot price regression residuals.

The final step in spot price simulation is to calibrate the resulting simulations to traded market data. For liquid locations, approaches include

—direct transformation of the physical spot simulations to a distribution with moments consistent with forward prices and, if desired, options. A combination of shift and dilation of the form $a(t) + b(t)p(t)$ are commonly used, although more refined methods can account for skew and can accommodate known bounds on prices.

—Alternatively, the parameters $[\alpha, \gamma, \vec{\lambda}]$ in (15) can be adjusted to minimize the difference between model prices and market trade prices.

The hierarchical approach is particularly suited to the inference of price distributions at illiquid nodes of the hierarchy. Given that the distribution at the parent node has been calibrated to market data, the distribution of prices at a child node can be computed by direct application of the calibrated model (15).

4. Application to Asset Valuation and Hedge Rating

In this section we will apply the econometric modeling framework described above to hedge a fictitious combined-cycle generator located in the PSEG Zone in PJM. This particular zone in the PJM market is reasonably liquid—forward prices are readily visible for at least several years of tenor. The asset that we will consider will deliver power at the FairLawn node—selected quite randomly for this example from many possible nodal delivery locations. A high-level summary of the asset is as follows:

—Generation capacity at 600 MW and a heat rate of 6.9,

—Additional generation capacity (duct firing) at 90 MW and a heat rate of 9.0,

—Power delivery price: Fairlawn 138 kV, and

—Fuel price: First 35% at TETM3, the remaining 65% at TZ6NY. These are two distinct and commonly traded northeastern natural gas hubs.

In addition to these attributes, a litany of other features including seasonal variation in the capacity and heat rates above, ramp rates, minimum run times, varying start costs, and emissions constraints were used to produce a realistic representation of an actual generator.

It is worth emphasizing that the underlying natural gas price is a blend of two distinct prices. This complexity is relatively simple to handle in the hierarchical setting.

The hedging period will be June 2018–May 2019. As done earlier, we will denote the random set of cash flows from this asset by Π. The hedge that we will analyze will be a revenue put—a currently popular method of supporting project development. This structure puts a floor (lower bound) on the revenue of a "synthetic asset," which, as before, we denote by $\tilde{\Pi}$. The revenue put has a strike of $K = \$60\text{m}$ so that the payoff to the hedger is $\max[K - \tilde{\Pi}, 0]$.

In what follows, the specification of $\tilde{\Pi}$ differs from the asset only in that the underlying prices are the PSEG spot price (as opposed to the nodal price), with constant specified adders to the peak and off-peak components.

The analysis required construction of the joint spot price distributions and calculation of the optimal dispatch Q in (1) for both the asset and the synthetic asset. The optimization method used to compute Q is not central to this paper, so we will not go into more detail here other than to say that a heuristic optimization method was used.

Figure 17 shows the synthetic payoff $\tilde{\Pi}$ versus asset payoff Π. These payoffs, while highly correlated, are clearly not identical. Figure 18 shows distributions of the asset cash flows and those hedged portfolio—namely, the sum of the asset and revenue put cash flows. The

FIGURE 17. Synthetic asset vs. asset payoffs.

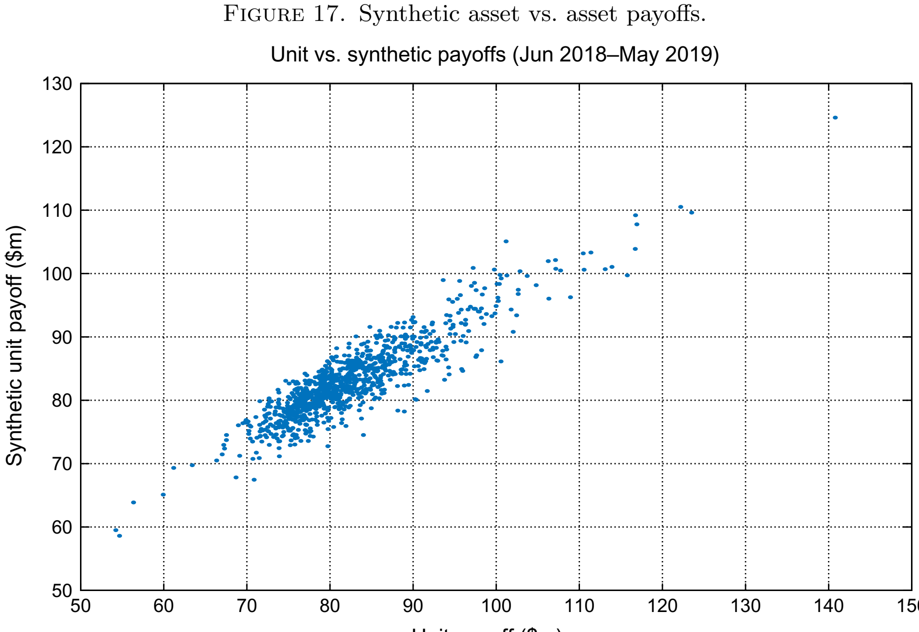

FIGURE 18. Payoff distributions pre and post hedge.

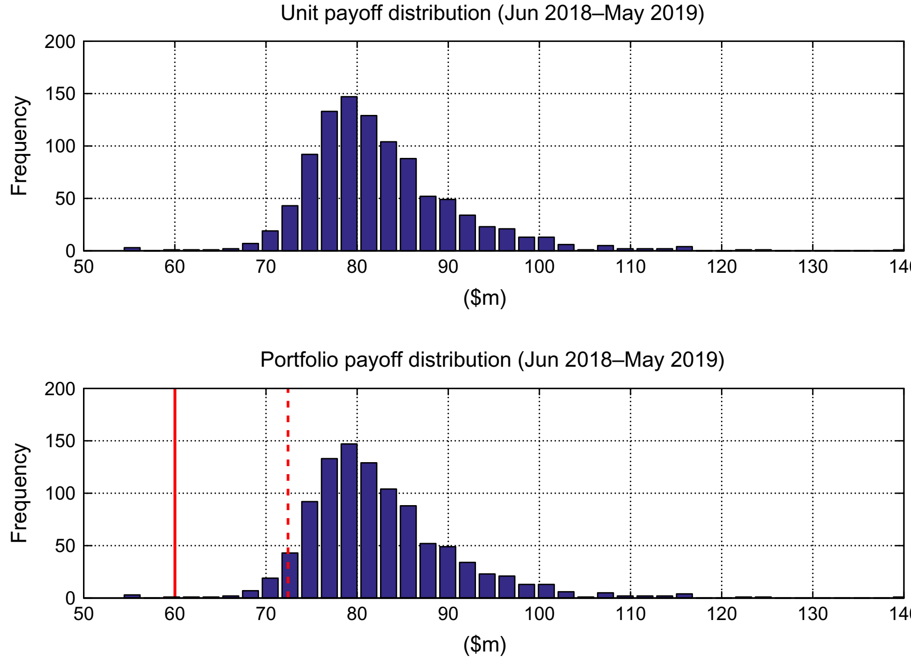

effect of the revenue put is minimal as few realizations of $\tilde{\Pi}$ are below the strike shown in solid red.

Perturbations to the current market are straightforward to effect in this framework. Lenders are (or, if not, should be) interested in how well a hedge would perform under a large and unforeseen change in markets conditions. Figure 19 shows similar results for a 20% drop

FIGURE 19. Payoff distributions pre and post hedge (low heat rate case).

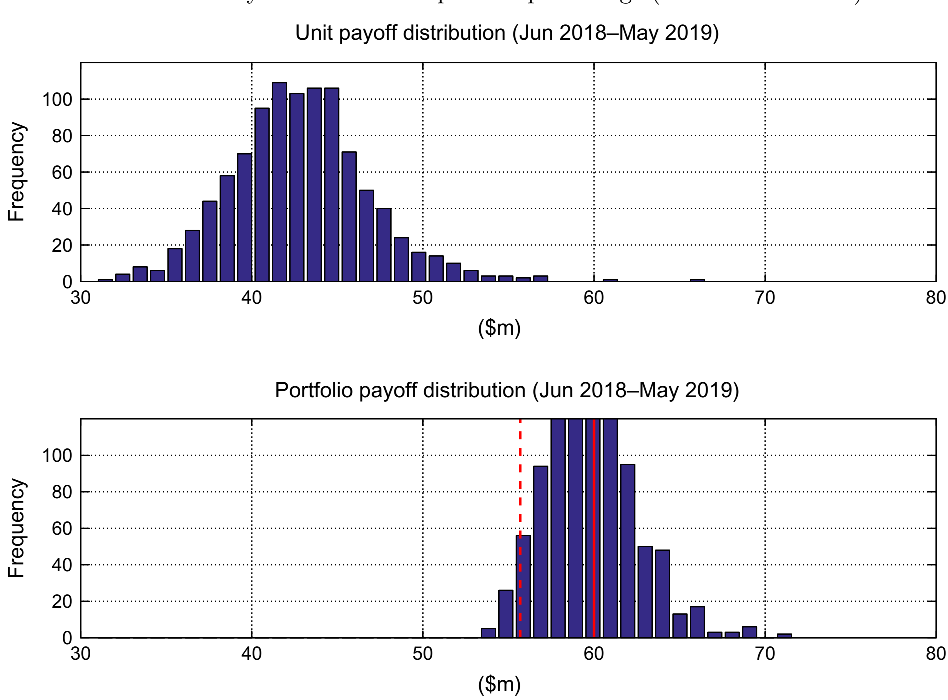

in the forward price of power versus natural gas—a large change that would substantially alter the solutions of the control problems and the resulting cash flow distributions. Here, one can see that the modeled impact of such a change is that the hedge does not in fact protect the financing structure at $60 million but at the fifth percentile (shown in dashed red) at approximately $55 million. This provides a rigorous and useful characterization of hedge performance that is difficult to achieve using purely reduced-form models.

5. Application to Variable Quantity Swaps

Variable quantity swaps were the second class of valuation problems introduced in the first section. Modifying notation to accommodate demand at the hourly timescale, a fixed price contract to serve a customer yields a cash flow in a given month m of the form

$$\Pi_m = \sum_{h\in m} L_h(p_f - p_h), \tag{19}$$

where L_h is the hourly demand, p_h the hourly spot price, and p_f the fixed price of the contract. The holder of such a position receives p_f for each megawatt-hour provided to the customer while procuring the commodity at the spot price p_h. The goal is to establish a fair-value price p_f and optimal hedges for (19).

Electricity forwards do not trade at the hourly timescale but for delivery buckets corresponding to peak and off-peak hours (peak being, for example, 7 A.M.–11 P.M. on business days). Without loss of generality, we will proceed by analyzing (19) for a single delivery bucket B in a given month.

Given a set of potential hedges $\vec{\mathcal{H}}$ with market prices $\vec{p}_{\mathcal{H}}$, assuming a specified fixed price p_f, the minimum variance portfolio is given by

$$\arg\min_{\vec{w}} \operatorname{var}\left[\sum_{d\in m}\left(p_f L_B(d) - \sum_{h\in B(d)} L_h p_h\right) + \vec{w}^{\dagger}(\vec{\mathcal{H}} - \vec{p}_{\mathcal{H}})\right]. \tag{20}$$

The optimal hedge is a function of p_f: $\vec{w}_*(p_f)$ and, therefore, must be solved jointly with p_f.

For a single hedging instrument, the solution is

$$p_f = -\frac{E[\widehat{LP}] - \operatorname{cov}[\widehat{LP}, H]E[H]/\operatorname{var}[H]}{E[\hat{L}] - \operatorname{cov}[L,H]E[H]/\operatorname{var}[H]} \tag{21}$$

and

$$w_* = \frac{\operatorname{cov}[\widehat{LP}, H] - p_f \operatorname{cov}[L,H]}{\operatorname{var}[H]}, \tag{22}$$

where

$$L \equiv \sum_{d\in m} L_B(d) \qquad \widehat{LP} \equiv \sum_{h\in B(d)} L_h p_h. \tag{23}$$

The result for w_* is the standard min-variance solution for the two payoff components. The first term is usually positive since high load is correlated with high price. If L and H are positively correlated, the second term reduces the short position.

To complete the analysis, the joint distribution of $[L, p, \widehat{LP}]$ or, equivalently,

$$\vec{X}_d \equiv [L_B(d), \mathcal{U}_B(d), p_B(d)] \tag{24}$$

is required, with

$$L_B(d) \equiv \sum_{h\in B(d)} L_h \qquad \mathcal{U}_B(d) \equiv \frac{\sum_{h\in B(d)} L_h p_h}{(\sum_{h\in B(d)} L_h) p_B(d)}. \tag{25}$$

FIGURE 20. Load vs. temperature with fit.

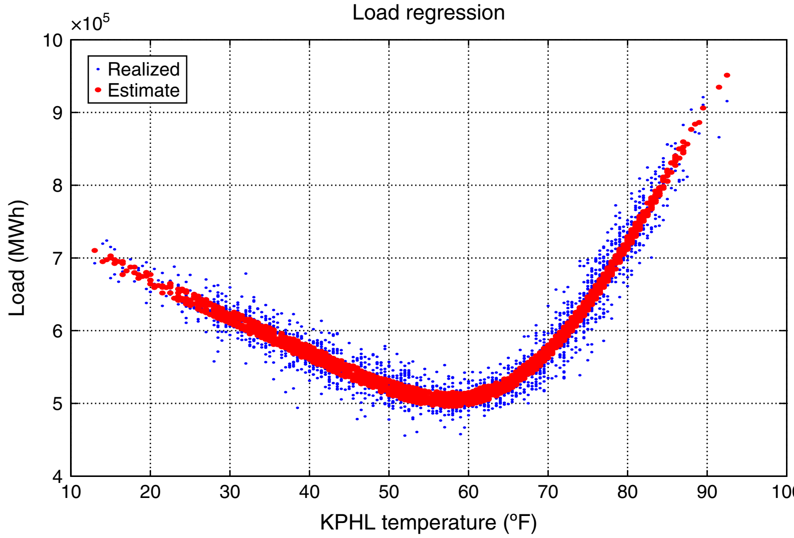

FIGURE 21. Load swap payoff vs. spot prices.

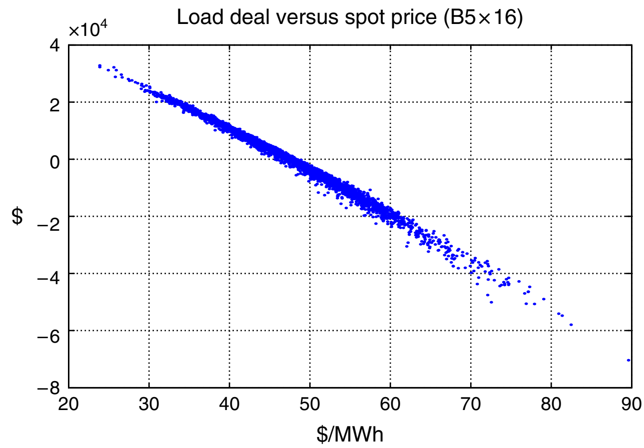

The following results are for pricing date April 21, 2016 for the delivery month July 2016. As for the generation example, we will take the spot index as the PSEG hourly day-ahead price. The demand L_h will be 10^{-4} of PJM Mid-Atlantic hourly load published by the ISO (Independent System Operator).[7]

Figure 20 shows historical PJM Mid-Atlantic load versus KPHL (Philadelphia) temperature for the peak bucket, in addition to a regression fit of the form

$$L_B(d) = \alpha + \beta d + \sum_{k=1}^{K_L} \theta^k(d) + \sigma_L(d)\epsilon_L(d). \tag{26}$$

A similar regression form was used for the $\mathcal{U}_B(d)$.

Finally, a choice of what hedges should be included in $\mathcal{H}$ must be made. Figure 21 shows the simulated payoff with p_f selected to be equal to the expected spot price. Being nearly linear in structure, a single forward contract is a reasonable choice for $\mathcal{H}$. Moreover, most

[7] PJM Mid-Atlantic is a large load base consisting of millions of consumers, hence the small fraction.

FIGURE 22. Hedge payoff vs. spot price.

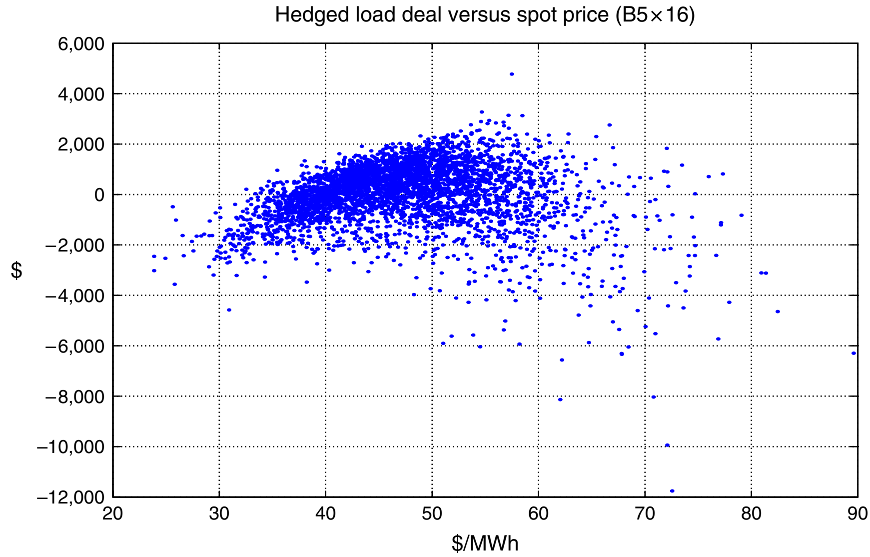

load serving entities use forward purchases as the primary hedging method. The results are as follows.

• The fair-value fixed price is $p_f = 49.816$ compared with the 5×16 forward price of 47.225. The fair-value price is roughly 5.5% higher than the standard forward price as a result of the adverse correlation between load and price at the hourly level.

• The optimal hedge is $w_* = 1{,}512$ MWh. This is approximately 11% higher than the expected load of 1,358 MWh, once again as a result of the adverse load/price correlation.

Figure 22 shows the hedge residuals versus power price. In practice, little more can be done about the residual risks without the use of heavily structured (and often expensive) hedges. This figure portrays the distribution of cash flows that can realistically be hedged by the holder of the short variable demand position.

The punchline is that this methodology not only establishes a fair-value price p_f and optimal hedge but also characterizes the residuals risk. As a final point, most variable demand portfolios contain similar transactions at many distinct delivery locations, a situation that is easily handled in the hierarchical framework.

6. Summary

Many energy structures are amenable to standard financial engineering methodologies. The application of risk-neutral valuation using reduced-form models is sensible in situations where the structures being valued depend on the prices at delivery locations, which support significant liquidity in both forwards and options.

Situations often arise, however, that push the requirements of an analytical framework into illiquid or untradeable risks at many price locations simultaneously. It is in such cases that econometric or structural models are particularly relevant, allowing price inference, optimal hedge construction, and a characterization of the distribution of resulting unhedgeable components of the portfolio.

The use of such frameworks is certainly far from being free of challenges. A lack of stationarity is a primary concern. And there have been many factors in play in recent years that result in nonstationary price dynamics. Fracking and shale gas production has dramatically changed the behavior of natural gas, and by extension power prices, across the United States. The addition of new sources of energy such as wind and solar generation have upended European power markets and are having an increasing impact on price dynamics in North America. As different energy commodities prices change in relative cost, the value of assets can change as substitution occurs.

Econometric models are partly insulated from nonstationarity by explicit calibration traded forward prices. However, if the factors affecting an illiquid pricing point are idiosyncratic, traded forwards are not of much use. This is where structural models, though more challenging to implement, have an advantage—the stylized representation of the physical market and price clearing afford the possibility of predicting changes in price dynamics as the physical system evolves.

The two examples discussed in detail were each in the context of electricity markets. This was for several reasons. First, electricity markets provide reliable and complete historical data. Second, structured hedges are commonly used in electricity markets to support project finance. Third, electricity markets support vibrant retail energy markets with many companies that routinely have literally thousands of variable demand positions in play and that are routinely having to grapple with the hedging problems illustrated above.

Finally, econometric and structural models are more well developed in the electricity markets than in other markets. The lack of storage of any appreciable size decouples price clearing at distinct times and simplifying model form significantly. Illiquid delivery locations and random quantities are, however, relevant to all energy commodities, and the development of viable structural models in storable commodities remains one of the most interesting areas of research in the energy commodities space.

References

[1] R. Aid, L. Campi, A. N. Huu, and N. Touzi. A structural risk-neutral model of electricity prices. *International Journal of Theoretical and Applied Finance* 12(7):925–947, 2009.

[2] M. Barlow. A diffusion model for electricity prices. *Mathematical Finance* 12(4):287–298, 2002.

[3] O. Barndorff-Nielsen, F. Benth, and A. Veraart. Modelling energy spot prices by Lévy semistationary processes. CREATES Research Paper 2010-18, University of Aarhus, Aarhus, Denmark, 2010.

[4] G. Barz and B. Johnson. Modeling the prices of commodities that are costly to store: The case of electricity. *Proceedings of the Chicago Risk Management Conference, Chicago*, 1998.

[5] M. Baxter and A. Rennie. *Financial Calculus: An Introduction to Derivatives Pricing*. Cambridge University Press, Cambridge, UK, 1996.

[6] F. E. Benth, J. Kallsen, and T. Meyer-Brandis. A non-Gaussian Ornstein–Uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance* 14(2):153–169, 2007.

[7] F. Black. The pricing of commodity contracts. *Journal of Financial Economics* 3(1):167–179, 1976.

[8] R. Carmona and M. Coulon. A survey of commodity markets and structural approaches to modeling electricity. F. Benth, V. Kholodnyi, and P. Laurence, eds. *Quantitative Energy Finance*. Springer-Verlag, New York, 41–83, 2013.

[9] R. Carmona and M. Ludkovski. Spot convenience yield models for the energy markets. G. Yin and Q. Zhang, eds. *Mathematics of Finance*, Contempoarary Mathematics, Vol. 351. American Mathematical Society, Providence, RI, 65–79, 2004.

[10] R. Carmona and M. Ludkovski. Valuation of energy storage: An optimal switching approach. *Quantitative Finance* 10(4):359–374, 2010.

[11] R. Carmona, M. Coulon, and D. Schwarz. Electricity price modeling and asset valuation: A multi-fuel structural approach. *Mathematics and Financial Economics* 7(2):167–202, 2013.

[12] Á. Cartea and M. Figueroa. Pricing in electricity markets: A mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* 12(4):313–335, 2005.

[13] L. Clewlow and C. Strickland. A multi-factor model for energy derivatives. Working paper, University of Technology Sydney, Ultimo, NSW, Australia, 1999.

[14] M. Coulon, W. B Powell, and R. Sircar. A model for hedging load and price risk in the Texas electricity market. *Energy Economics* 40(November):976–988, 2013.

[15] S. Deng. Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes. Technical report, University of California Energy Institute, Berkeley, 2000.

[16] A. Eydeland and K. Wolyniec. *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*. John Wiley & Sons, Hoboken, NJ, 2003.

[17] H. Geman. *Commodities and Commodity Derivatives: Modeling and Pricing for Agricultural, Metals and Energy*. John Wiley & Sons, Chichester, UK, 2005.

[18] H. Geman. Pure jump Lévy processes for asset price modeling. *Journal of Banking and Finance*, 32:2520–2529, 2008.

[19] H. Geman and A. Roncoroni. Understanding the fine structure of electricity prices. *Journal of Business* 79(3):1225–1261, 2006.

[20] R. Gibson and E. S. Schwartz. Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance* 45(3):959–976, 1990.

[21] S. Grine and P. Diko. Multi-layer model of correlated energy prices. *Journal of Computational and Applied Mathematics* 233(10):2590–2610, 2010.

[22] S. Hikspoors and S. Jaimungal. Asymptotic pricing of commodity derivatives using stochastic volatility spot models. *Applied Mathematical Finance* 15(5–6):449–477, 2008.

[23] V. Kaminski. The challenge of pricing and risk managing electricity derivatives. R. Jameson, ed. *The US Power Market*. Risk Publications, London, 149–171, 1997.

[24] V. Kholodnyi. Modeling power forward prices for power with spikes: A non-Markovian approach. *Nonlinear Analysis: Theory, Methods and Applications* 63(5–7):958–965, 2005.

[25] C. Pirrong. *Commodity Price Dynamics*. Cambridge University Press, Cambridge, UK, 2012.

[26] D. Ribeiro and S. Hodges. A two-factor model for commodity prices and futures valuation. *EFMA 2004 Annual Meetings Paper, Basel*, 2004.

[27] P. Skantze, A. Gubina, and M. Ilic. Bid-based stochastic model for electricity pricing. MIT E-Lab Report MIT-EL 00-004, Massachusetts Institute of Technology, Cambridge, 2000.

[28] P. Skantze, M. Ilic, and J. Chapman. Stochastic modeling of electric power prices in a multi-market environment. *Proceedings of the IEEE Power Engineering Society Winter Meeting 2000*, Vol. 2. IEEE, Piscataway, NJ, 1109–1114, 2000.

[29] G. Swindle. *Valuation and Risk Management in Energy Markets*. Cambridge University Press, New York, 2014.

[30] A. B. Trolle and E. S. Schwartz. Unspanned stochastic volatility and the pricing of commodity derivatives. *Review of Financial Studies* 22(11):4423–4461, 2009.