Shaping and Hedging Renewable Power Purchase Agreements

Brock Mosovsky, Ph.D.
Director of Operations and Analytics, cQuant.io

Lance Titus
Managing Director, Uniper Global Commodities; and Member of the J.P. Morgan Center for Commodities' (JPMCC's) Research Council at the University of Colorado Denver Business School

Mr. Lance Titus (right), Managing Director, Uniper Global Commodities, presenting during the panel on "Emissions Trading" during the JPMCC's Research Council meeting on September 30, 2016. Mr. Titus is a member of both the JPMCC's Research Council and its Advisory Council and also serves on the Editorial Advisory Board of the *Global Commodities Applied Research Digest*. Dr. Daniel Kaffine (left), Ph.D., Professor of Economics at the University of Colorado Boulder, also participated in the "Emissions Trading" panel.

Renewable power purchase agreements (PPAs) have steadily increased in popularity over the last decade. They have enabled hundreds of megawatts of renewable energy development and have played important roles in many corporate and utility-led sustainability programs. As renewable PPAs have been accepted by a broadening range of market participants and as the intermittent nature of renewable generation has become better understood, PPAs themselves have increased in sophistication. Today, many PPAs include provisions designed to protect against uncertainty in renewable energy generation or volatile electricity market prices. Collectively, these protections fall under the umbrella of “shaping and hedging” and aim to provide one counterparty or another with increased certainty around future generation or cash flows.

This article is the second in a two-part series on renewable PPA valuation and risk assessment published in the J.P. Morgan Center for Commodities’ (JPMCC’s) *Global Commodities Applied Research Digest* (*GCARD*).[1] Herein, we outline methodologies for shaping and hedging renewable PPAs, and we discuss the benefits of each of these strategies from both the buyer and seller perspectives. We build up PPA value across a range of shaping scenarios of increasing granularity, identifying the incremental value of each refinement in shape. We also present a methodology for deriving optimal hedge ratios that can be used to enact a hedging program that minimizes risk to the buyer or seller and that is custom-tailored to a particular renewable facility or PPA.

Shaping Renewable Energy

In a shaped PPA, the seller guarantees the buyer a fixed generation shape—a predetermined quantity of energy delivered over a predetermined period of time. In exchange, the buyer guarantees that the seller will be compensated at the PPA price for all energy delivered under the contract or financially settled as a contract for differences (“CFD”).[2] The shape guarantee may apply at the annual, seasonal, monthly, or even hourly level and, depending on the granularity of the shape profile, has the effect of removing some or all uncertainty in renewable generation (generation risk) from the buyer’s position in the contract.

In addition to removing risk from the buyer’s position in the PPA, shaping also has other benefits to the buyer. It can help align energy contracted under a PPA with the shape of the buyer’s load (his native short position), giving the buyer greater confidence in managing his residual load position. A corporate buyer’s load may be relatively flat compared to the highly variable generation produced by a wind or solar farm, and a shaped PPA can help to align supply and demand in a more predictable way than a unit-contingent or “as produced” PPA. In cases where a buyer’s electricity tariff is directly related to wholesale market prices, the improved alignment with the buyer’s load may allow the PPA to function as a better financial hedge against the buyer’s native short physical position, providing protection against future electricity price fluctuations. Finally, shaping a PPA can better align contracted generation volumes with standard over-the-counter financial products, allowing the buyer to directly lock-in future value through hedging.

Such benefits of a shaped PPA do typically come at a cost, however, and sellers will demand a premium to assume the generation risk on behalf of the buyer. In any contracting scenario, the question becomes: how much of a premium is reasonable? The answer is highly dependent on the contract terms, the location of the facility, and the real-time dynamics of the renewable resource at the point of generation and the electricity prices at the settlement point. Each PPA has its own unique profile of value and risk and requires rigorous case-by-case analysis to properly understand the contractual implications to the buyer and seller.

Dr. Brock Mosovsky, Ph.D. (left), Director of Operations and Analytics, cQuant.io, with his colleague, Mr. David Leevan (right), Managing Director, cQuant.io, during the JPMCC's 2nd International Commodities Symposium, which was held at the University of Colorado Denver Business School on August 14 through August 15, 2018.

Shaped PPA Settlement Amounts

For all examples in this article, we assume that shapes guarantee a certain amount of generation in each hour (as opposed to each month, season, or year), though guaranteed volumes may vary hour-to-hour and month-to-month. This structure is sometimes known as an "8760 profile" since there are 8760 hours in a typical (non-leap) year and the shape guarantees a specific amount of energy to be delivered in each hour. As in the first article of the series, we also assume contracts are for virtual PPAs that settle financially each month. Under these assumptions, the buyer's settlement amount, $\boldsymbol{A_{buy}}(\boldsymbol{h})$, in any particular hour, h, for a shaped PPA is given by the equation,

$$\boldsymbol{A_{buy}}(\boldsymbol{h}) = G(h)[\boldsymbol{p}(\boldsymbol{h}) - K], \#(1)$$

where $G(h)$ is the guaranteed contracted generation amount in hour h, $\boldsymbol{p}(\boldsymbol{h})$ is the variable market price of electricity in hour h, and K is the fixed PPA price in dollars per megawatt-hour ($/MWh). Here we have used the convention that terms in bold font denote quantities that are uncertain in each hour over the contract horizon; these are the terms that impart risk to the PPA. From equation 1, it is easy to see that the buyer assumes no generation risk. That is, in any given hour, the only uncertain quantity is the market price, $\boldsymbol{p(h)}$.

The hourly settlement amount, $\boldsymbol{A_{buy}(h)}$, is the amount the buyer pays or receives in a given hour under the PPA. It may be either positive or negative, depending on the market price of electricity, $\boldsymbol{p(h)}$, in relation to the contract price, K. When $\boldsymbol{p(h)}$ is greater than the contract price, the buyer receives a payment; when $\boldsymbol{p(h)}$ is less than the contract price, the buyer makes a payment. Put another way, equation 1 is the buyer's hourly cash flow under the contract. It is equivalent to saying that the buyer pays the contracted price, K, for the guaranteed generation, $G(h)$, and also receives a payment in the amount of the real-time market price valuation of that generation.

Typically, PPAs settle monthly, which means that no cash actually changes hands until the end of the month. The monthly settlement amount is simply the sum of the hourly settlement amounts over all hours of the month. For the buyer, this is,

$$\boldsymbol{A_{buy}} = \sum_{h\ in\ month} \boldsymbol{A_{buy}(h)}. \#(2)$$

The seller's risk profile differs significantly from the buyer's risk profile. In a shaped PPA, the seller explicitly assumes the generation risk in the contract by guaranteeing some production level to the buyer. In practice, this guarantee is usually made at the P99 level, or the production profile that will be met with 99% statistical confidence; the P99 profile is often used to size debt service coverage ratios for financing new renewable projects. When actual facility production is below the shaped level, the seller must purchase energy from the market to make up the difference; when actual production is above the shaped level, the seller may liquidate the residual energy into the market at the prevailing real-time price. This is in contrast to a unit-contingent PPA where no guarantee on generation is made and the buyer simply accepts all or a prorated share of the energy produced by the renewable facility in each hour. Regardless of generation level, the seller of a shaped PPA is guaranteed to receive the contracted price, K, for each unit of energy covered under the shape. As such, the seller's settlement amount, $\boldsymbol{A_{sell}(h)}$**,** in a particular hour, h, for a shaped PPA is given by the equation,

$$\boldsymbol{A_{sell}(h)} = G(h)K + [\boldsymbol{g(h)} - G(h)]\boldsymbol{p(h)}, \#(3)$$

where $\boldsymbol{g(h)}$ is the actual generation produced by the facility in hour h and all other terms are as defined in equation 1. This equation states that the buyer's hourly settlement amount is the guaranteed generation (the contracted shape) valued at the contracted price plus the difference between actual and contracted production valued at the prevailing real-time market price.

Rearranging terms in equation 3 helps to more clearly isolate the elemental components of risk for the seller:

$$\boldsymbol{A_{sell}(h)} = G(h)K + \boldsymbol{g(h)p(h)} - G(h)\boldsymbol{p(h)}. \#(4)$$

From this equation, we see that the first term, $G(h)K$, is entirely determined at contract signing and contains no uncertain quantities. This represents the guaranteed payment from the buyer for the contracted energy shape. The second term, $\boldsymbol{g}(\boldsymbol{h})\boldsymbol{p}(\boldsymbol{h})$, is the product of two uncertain quantities in each hour: the realized hourly generation and the real-time market price. The fact that these two uncertain quantities are multiplied together has important implications for the seller's risk; effectively, the generation and price risk can have a magnifying effect on each other. We discuss this magnification of risk in more detail below in the section on hedging. The final term in equation 4, $G(h)\boldsymbol{p}(\boldsymbol{h})$, contains only price risk and is identical to the buyer's price risk in equation 1 up to a difference in sign.

Example PPAs – Wind and Solar in Texas

In order to demonstrate the practical implications of the shaping equations above, the value of shaping contracts at varying levels of granularity, and the optimal ways in which shaped contracts can be hedged, we examine several virtual PPAs settled against Electric Reliability Council of Texas (ERCOT) North Hub real-time prices. The PPAs are based on hypothetical wind and solar farms in central Texas located approximately 150 miles northwest of San Antonio. The solar farm is assumed to be a 10 MW DC fixed array with a panel tilt angle optimized for maximum annual energy production. The wind farm is assumed to be composed of Siemens SWT-2.3-108 wind turbines, and the full wind farm capacity is scaled so that its expected annual energy production is identical to that of the solar farm. This scaling eliminates any volumetric bias in aggregate between the two renewable resources and their valuations. Since we also use a consistent set of simulated market prices[3] to value the solar and wind PPAs, any differences in value and risk are the direct result of differences in the timing of generation and its alignment with both market prices and contract specifications. We assume a 5-year contract term beginning in January of 2019 and running through December of 2023.

A primary driver of PPA value is the alignment between the shape of generation and the shape of prices at both the monthly and hourly levels. The fair market value of a given PPA is essentially the expected generation-weighted average market value of energy to be produced under the contract. As such, generation during hours when market prices are above average will increase PPA value while generation during hours when market prices are below average will decrease value. In either case, higher generation levels in a given hour will produce a larger effect on PPA value for that hour. That is, when all other variables are held fixed, a generation shape that is highly coincident with the shape of market prices will have more value than one that is misaligned with market prices.

Figure 1 on the next page shows the alignment of median hourly solar and wind generation shapes (known as the P50 shape) with ERCOT North real-time prices in each calendar month. The plots show a stark contrast between the coincidence of the two generation types with market prices. Summer ERCOT North electricity prices show a very strong hourly shape with evening prices being more than twice the value of early morning prices at the median from June through September. Solar generation takes advantage of this strong mid-day rise in prices with peak power output occurring just a few hours before the evening peak price. Despite the solar peak occurring slightly before the evening price peak, there is generally a strong correspondence between elevated solar generation and elevated price, resulting in enhanced overall value for solar PPAs.

Figure 1
Median (P50) Hourly Generation and Price Shapes

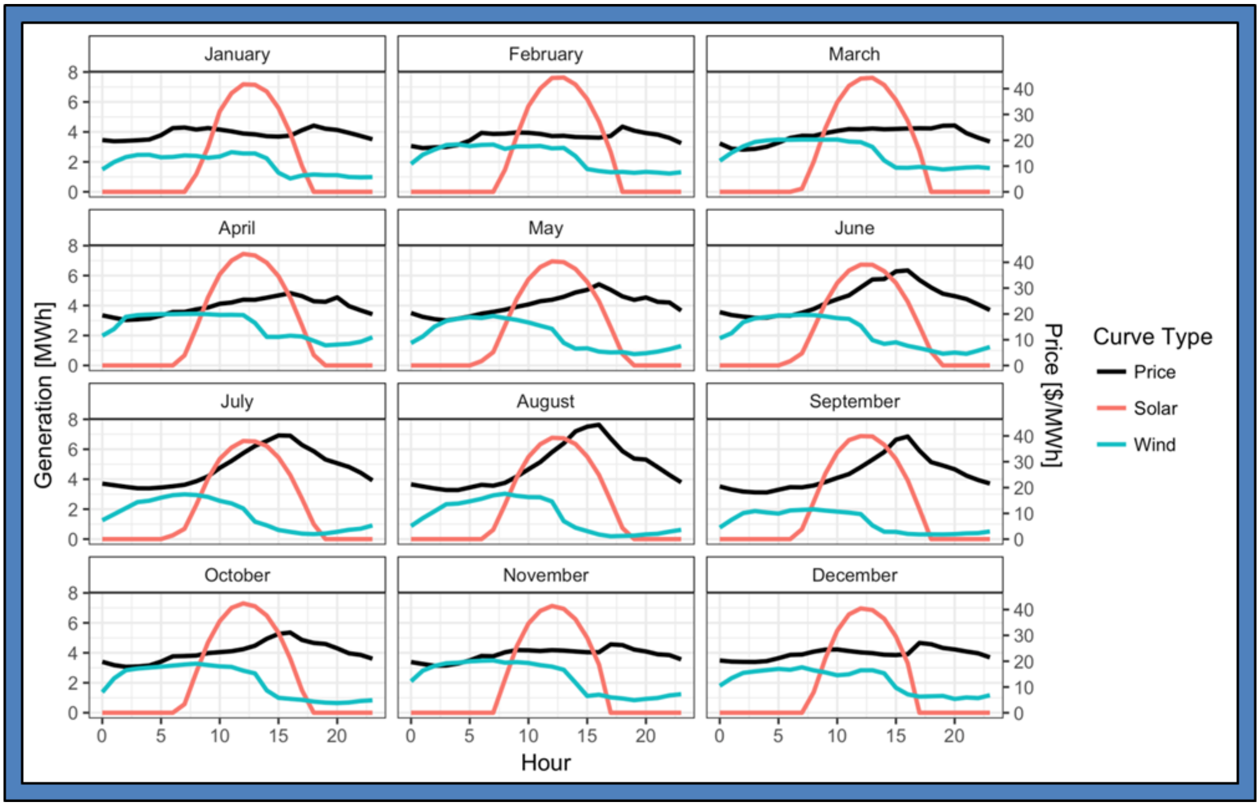

Notes: Generation curves reflect hourly median simulated generation from 2007 through 2012 based on historical wind speed and solar radiation data obtained through the National Solar Radiation Database and the NREL Wind Integration National Dataset. Price curves reflect hourly median ERCOT North real-time spot prices over the period from January 2014 through June 2018.

Source of image: cQuant.io ReAssure PPA®.

Wind generation, on the other hand, peaks in the early morning hours and tends to ramp down just as prices begin to rise in the early afternoon. The median wind generation shape virtually mirrors the price shape in May through October; this assigns greater weight to hours with lower market prices and tends to drag PPA value downward. The summer misalignment between wind generation and price is only partially mitigated in the winter when elevated prices in the early morning do coincide with high wind generation. However, the winter evening hours show the opposite trend with wind generation falling off just before the evening peak in price.

While the plots in Figure 1 help to paint a picture of the value of shaping the PPA relative to intra-day patterns in generation and prices throughout the year, they do not provide the full story. Equally important to price shape is the absolute price level in each month. One indicator of the market's expectation of the future price of electricity is the "forward curve," or the set of forward or futures contracts traded today for delivery of energy at some point in the future. The historical shape analysis can be combined with this current market view of future electricity prices to provide a "mark-to-market" fair market valuation of the shaped energy over the life of the PPA. Another important aspect missing from the P50 shapes discussed above is an indication of the level of uncertainty in generation and price in each month. The way in which this uncertainty is actualized over the life of a PPA can have a significant effect on overall value. As such, it is important to properly incorporate uncertainty in any assessment of contract value or risk.

In the first article in the series, we provided details on the mechanics of combining historical spot price analysis with current forward curves to inform PPA fair market value. We also discussed how a simulation-based approach can provide an understanding of uncertainty in both generation and price. In this article, we focus on the end results of the simulation-based analysis for determining PPA value and risk. In particular, we investigate various different granularities of shaped generation and the value and risk that different shape profiles impart to the contract.

Figure 2 below presents a breakdown of various different generation shape components that affect the fair market value of a renewable PPA. In order to isolate the true value of each shape component, both generation and market prices are considered at the same level of granularity when computing overall contract value. Moving from left to right along the x-axis in the charts, each successive shape scenario uses an increasingly granular methodology for aligning generation with price to value the PPA, and the incremental change in value at each successive step represents the fair market value for that particular shape component.

Figure 2

Average Value of Energy Generated for an Example Solar and Wind PPA in Texas under Various Shaping and Valuation Granularities

Notes: Each bar shows the incremental value to the PPA for a particular shape scenario. Simulations of market prices incorporated quotes for ERCOT North Hub real-time electricity futures contracts obtained from cmegroup.com as of June 22, 2018.

Source of image: cQuant.io ReAssure PPA®.

It is important to note that the value and risk of the various shape components can vary widely from location to location and as electricity price dynamics and grid topology evolve over time. To ensure all important locational parameters for a given renewable generation facility are captured and market dynamics are up-to-date, such an analysis should be performed periodically for each location in question. Additionally, the present analysis assumes that generation and market price dynamics are not directly coupled by any structural mechanism and the uncertainty in generation is assumed independent from the uncertainty in market prices. While this assumption is valid in markets with low to modest levels of renewable energy, significant structural relationships between renewable generation and market prices may develop in markets with deeper renewable penetration. In these cases, large concentrations of intermittent generation may have a material impact on electricity prices during periods of high production. A full analysis of the implications of such structural relationships, sometimes referred to as "renewable penetration risk," is outside the scope of the present discussion.

Baseload Valuation

The Baseload scenario in Figure 2 assumes the same volume of energy is guaranteed over every hour of the year and uses a single average annual price to value this energy. This is the coarsest shape one could achieve for the energy under a PPA and also the coarsest approach one could take to value that energy. The hourly shaped quantity is simply the expected total annual energy for a single year divided by the number of hours in that year, and the fair market value is the average around-the-clock (ATC or 7x24) forward contract price over the year, weighted by the number of hours in each month. As Figure 1 clearly shows, in reality there is significant seasonal and hourly variation in both generation and price. Nonetheless, the Baseload valuation approach does capture a majority of the PPA's value by accounting for the average hourly generation and the average price that generation would receive in the market.

Flat Monthly Valuation

Recognizing that the Baseload approach provides far too coarse a view of value, the Flat Monthly scenario shapes both generation and price by calendar month. This valuation approach incorporates seasonal effects by using the monthly forward contract prices for electricity and weighting these prices by the expected renewable generation in each corresponding month. That is, the flat monthly shape guarantees delivery of a fixed quantity of energy in each hour of a given month, though this quantity does vary month-to-month to account for seasonal effects. This shaping and valuation approach results in a small increase in overall PPA value for solar and a slight decrease in value for wind. The overall direction of the incremental change compared to the Baseload scenario is determined by the seasonal alignment of generation and price. This impact is positive for solar and negative for wind generation, consistent with seasonal production characteristics and the term structure of prices for ERCOT North Hub in the summer months.

Figure 3 on the next page provides a more detailed view of the seasonal fluctuations in expected generation and price for the year 2019. Seasonal generation and price alignment is marginally coincident for solar, corresponding to a slight incremental increase in seasonally-weighted value compared to the Baseload scenario, as seen in Figure 2. For wind, seasonal generation and prices are somewhat anti-coincident, yielding a small loss in value over the Baseload shape scenario. In either case, the change in value from the baseload case is relatively small, indicating that seasonal effects are "averaged out" relatively well when taking a coarser annual view. Again, the result presented here is very location-specific; seasonal generation and price shapes may have different or more significant effects on PPA fair market value in other areas of the country due to differing weather and price dynamics.

Figure 3
2019 Monthly Expected Generation and Around-the-Clock (ATC) Forward Contract Prices

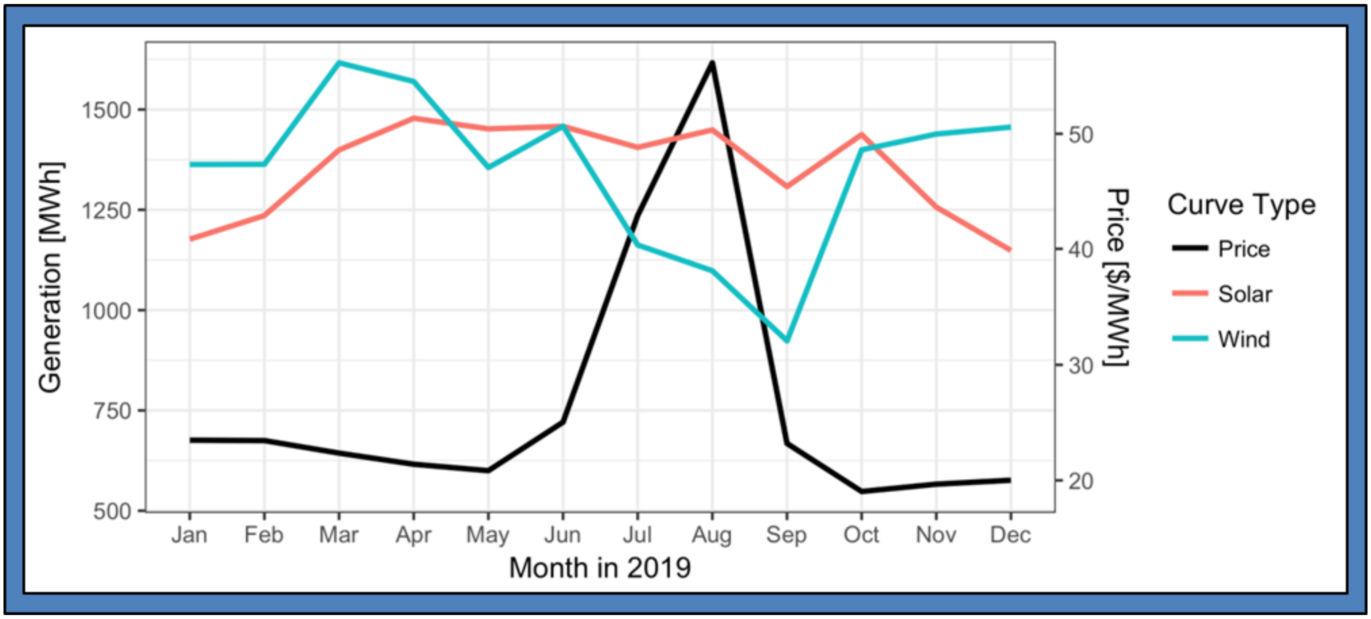

Note: Solar and wind display somewhat reversed seasonal trends, with solar peaking in the spring and summer and wind peaking in the winter and early spring. The solar facility represented in the figure is assumed to be a fixed array with a tilt angle optimized to maximize expected total annual energy production.

P50 Hourly Valuation by Month

Further increasing the granularity of the generation and price shaping, the P50 Hourly scenario values the PPA by aligning the median hourly generation and price shapes in each month of the year.[4] In addition to capturing seasonal changes in monthly aggregate generation and price levels, this methodology also captures the important hourly shapes illustrated in Figure 1, including the variation in hourly shape from month to month. Figure 2 shows that incorporating the hourly shapes into the valuation significantly affects fair market value for both solar and wind PPAs. Essentially, hourly shaping provides a more granular view of the intra-day alignment of generation with price. This results in a significant refinement in the weighting assigned to the market value of shaped generation in each hour compared to the Baseload and Flat Monthly scenarios.

The P50 Hourly valuation scenario differs fundamentally from the Baseload and Flat Monthly scenarios, which both include a guaranteed shaped amount of energy that does not fluctuate throughout a given day. A flat hourly energy shape is clearly misaligned with solar generation, which will always be zero during the nighttime hours when the sun is not shining. Moreover, the very hours in which solar does not generate tend to exhibit the lowest electricity prices, so overall PPA value will tend to be reduced by any shape that guarantees delivery over these hours.

Using the P50 Hourly profiles to shape the contract yields a guaranteed delivery profile that is much more consistent with actual solar production and emphasizes the natural alignment of solar generation with higher market prices during the peak hours of the day. This improved alignment in generation and price contributes an incremental increase in solar PPA value of more than 10% over the Flat Monthly valuation approach, as seen in Figure 2.

In contrast to solar, wind tends to show elevated generation during hours of the day when prices are low, as discussed above and shown in Figure 1. This anti-coincidence with price contributes an incremental decrease in wind PPA value for the P50 Hourly scenario compared with the Flat Monthly scenario. It is important that both the buyer and seller understand the negative incremental value of the hourly wind shaping with respect to price; over the life of a five-year 10 MW PPA, the almost $1.40/MWh reduction in the P50 Hourly shaped value of wind energy seen in Figure 2 translates to over $500 thousand in total value loss relative to the Flat Monthly scenario.

Full Stochastic Valuation and the Value of Uncertainty

The final shape scenario shown in Figure 2 uses a Full Stochastic valuation to simulate generation and prices and determine the true value of uncertainty around the shaped hourly P50 profiles. This value of uncertainty is akin to “extrinsic value” in the context of options theory or thermal generation asset valuation. The Full Stochastic valuation scenario represents an unshaped PPA, or what is often referred to as “unit contingent” or “as produced” in contracting terminology. That is, no guarantee is made by the seller about the energy produced under the PPA on any timescale.[5] In this case, the buyer assumes both the generation and price risk and takes whatever energy is produced by his full or prorated share of the facility. In some cases, the nature of the distributions of price and generation is such that the added uncertainty may actually contribute significant value to the contract, on average. This is particularly true when distributions of either generation or price (or both) are strongly positively skewed. Indeed, for both the solar and wind facilities analyzed in the present example, the uncertainty in generation and price actually comprises the second-largest component of total value for the contract, as seen in the right-most column of the solar and wind plots in Figure 2.

The ERCOT electricity market, in particular, is notorious for large price spikes during periods when the grid is strained due to extreme hot or cold weather (as covered in O’Neill (2017)), or when unexpected outages force significant portions of the generation stack out of operation. These price spikes result from “scarcity pricing,” a market mechanism whereby a large premium may be paid for energy generated during periods of low resource availability.[6] The result of this dynamic is that the distribution of historical ERCOT real-time prices is anything but normal, with an extreme positive skew and large excess kurtosis. For the historical ERCOT North Hub real-time hourly dataset used in the present analysis, the raw data showed a skewness of 33.5 and an excess kurtosis of almost 1400, whereas a normal distribution would have values of zero in both cases. The main point here is that the ERCOT North Hub shows an extremely significant directional shape in the uncertainty around expected prices, and this shape has a material effect on the valuation of energy flows within that market. The Full Stochastic valuation scenario accounts for this uncertainty shape, placing a dollar-per-megawatt-hour value on the combined uncertainty from both generation and market electricity prices for the contract.

In the case of the solar PPA in our example, the uncertainty in generation and price adds an additional $3.41/MWh to the value of the PPA above the P50 Hourly scenario. For wind, the incremental value is even greater at $6.39/MWh. In both cases, the mechanism driving the additional value is that both wind and solar may be generating during periods when the grid is strained and prices spike to extremely high levels. The difference between the value of uncertainty for solar and wind is primarily driven by the amount of uncertainty in the generation itself, as well as its alignment with price uncertainty. Wind generation tends to be more uncertain hour-to-hour than solar, and uncertainty peaks during summer months when market prices are also highly volatile and likely to undergo very high spikes. Even though expected wind generation dips in the summer months, the variability around this expected generation aligns well with the potential for very high prices; the net result is a significant positive $/MWh value of the uncertainty around expected generation and price.

Solar shares in this value, though to a lesser extent, since its hour-to-hour generation is less variable than wind and its seasonal profile already expects high generation in the summer months. Essentially, there is less upside potential for summer solar generation than there is for summer wind generation, and the occasional alignment in generation upside with large price spikes results in a disproportionate increase in value relative to the mean. The bars for the Full Stochastic scenario in Figure 2 illustrate the contribution of uncertainty to overall PPA value. For solar, the value of uncertainty makes up about 11% of total PPA value, while for wind it is an even higher proportion at roughly 21%. As the plots show, a proper understanding of the value of uncertainty in renewable PPAs can be absolutely critical to an accurate valuation of the contract. In the case of a five-year 10 MW wind PPA, the $6.39/MWh value of uncertainty observed in this example would translate to over $2.5 million in expected generation value over the contract lifetime.

While the analysis presented here does indicate a significant value contribution from uncertainty around the P50 shape, it should be noted that the unit-contingent or “as produced” contract type modeled by the Full Stochastic valuation scenario also involves more risk to the buyer than the various shaped deals. This is because the buyer accepts the generation risk in a unit-contingent contract whereas he is guaranteed a generation shape in the others. That is, while the added uncertainty may contribute additional value at the mean, the buyer is taking on a greater risk to access that value and must accept greater downside potential as well. Moreover, the fact that there is no guarantee of generation volumes means the unit-contingent contract is more difficult to hedge using standard financial products, as discussed in the following sections. In general, a PPA buyer should be aware of the full spectrum of risk before entering into any PPA, and he should have a plan in place to monitor and mitigate some of the long-term risks associated with the contract as market dynamics evolve.

Hedging Renewable PPAs

In the previous section we identified how various different shape granularities—Baseload, Flat Monthly, P50 Hourly, and Full Stochastic—contribute value to wind and solar PPAs in central Texas. However, restricting the granularity of the prices used to value a shaped PPA, as done above, is a somewhat artificial construct useful primarily for its ability to isolate the individual value contribution of each shape component. In practice one should always use a full stochastic simulation-based approach to value a PPA regardless of its contracted shape. That is, even if the PPA were contracted to include a baseload guaranteed shape of energy, a proper valuation should incorporate a robust understanding of the uncertainty in generation and price around the contracted quantity. The amount of residual (or deficit) energy between the contracted shape and the realized generation volume contributes risk to the seller's position in a shaped PPA and should be valued within this context.

In this section, we shift our focus from value to risk, examining cash flow uncertainty around renewable PPAs from both the buyer and seller's perspectives and exploring some strategies to mitigate this uncertainty. To simplify discussion, we consider contracts for two of the shape scenarios discussed above: one baseload shape for the same fixed number of MWh in every hour of the year and one shaped by the P50 hourly generation profile in each calendar month. We simulate settlement of these two contracts for the month of August 2019, consistent with the price simulations used to value the various shape scenarios above, and we demonstrate how a stochastic analysis can be used by both PPA buyers and sellers to design optimal hedges custom-tailored to a particular contract. We have selected the month of August because it typically contains a large amount of uncertainty in the ERCOT electricity market; however, the analysis presented here can be generalized and performed for all months of a PPA's horizon to design a hedging program that covers the entire contract.

Settlement Amounts Including Hedge Payoffs

As discussed above, the buyer and seller bear fundamentally different risks when entering into a shaped PPA. Because of the shape guarantee, the buyer is exposed only to price risk while the seller is exposed to both price and generation risk. As such, hedging programs should be designed with a thorough understanding of each counterparty's specific risk profile, appetite for risk, and primary objectives for entering into the PPA. While a comprehensive discussion of hedge program design is outside the scope of this article, we present examples that make use of around-the-clock (ATC) forward contract hedges enacted at the time of PPA signing.[7] The methodology we develop below can also be used to hedge PPAs using more advanced financial instruments and/or structured transactions and to layer risk-reduction programs onto an existing PPA after contract signing.

In the broadest sense, a hedge is a physical or financial position taken to reduce uncertainty in portfolio returns or cash flows. An effectively hedged portfolio contains multiple components where losses in one component are at least partially offset by gains in another. The general idea of offsetting returns is related to the concept of diversification from modern portfolio theory. Because of the availability of a number of financial derivative products designed specifically for the purpose of hedging energy positions, it is often easier to achieve a diversification effect in energy portfolios than it may be in portfolios of stocks. To hedge a given position, one may simply take an opposing position in a derivative contract that is similar in delivery volume, timing, and location.[8] The most common of these energy derivatives are swaps, forward contracts, futures contracts, and European options. Each of these, while they differ slightly in their exact contract specifications and settlement procedures, entitles the holder to buy or sell a pre-specified quantity of energy or other commodity at a pre-specified price at some point in the future.

For someone with an expected long future exposure in the spot market, taking a short position in a swap for the same commodity will replace exposure to the variable spot price with exposure to the fixed contract price at delivery. That is, it will "swap" a fixed price exposure for the native variable price exposure. This has the effect of reducing variability in the cash flows during the delivery period, and the magnitude of the reduction is proportional to the magnitude of the volume contracted in the swap. Essentially, by selling the swap, the seller is "locking in" the contracted price of the swap instead of remaining exposed to the uncertain future price in the spot market.

The Buyer's Position

From the above explanation of hedge implementation, it is clear that the variability in hedged portfolio returns is dependent on the quantity hedged. In the context of renewable PPAs, we define the hedge ratio, r, to be the ratio of the volume contracted in a hedge position to the expected volume in the native PPA position. For example, if a PPA is expected to deliver 1000 MWh in a given month and the buyer sells 300 MWh in a forward contract, this would correspond to a hedge ratio of 0.3. Using this definition, suppose a buyer of a baseload PPA decides to hedge his settlement amounts in a particular month by selling a swap at a price F for some ratio r of his guaranteed hourly generation quantity. We assume the price of the swap is equal to the average simulated price of power in the delivery month; this corresponds to the PPA and hedge valuations being effectively marked-to-market against the same forward curve. In practice, this would occur if the hedges were enacted at the time of PPA signing.

Because the contracted baseload generation from the PPA does not vary by hour, we modify the notation used in equation 1 and denote the fixed hourly generation quantity by G, dropping the dependence on time. Expanding the right-hand-side of equation 1 and including the payoff from the hedge, the hourly hedged settlement amounts for the buyer of a baseload PPA are given by

$$\begin{aligned} \boldsymbol{A_{buy}}(\boldsymbol{h}) &= G\boldsymbol{p}(\boldsymbol{h}) - GK - rG[\boldsymbol{p}(\boldsymbol{h}) - F]\#(5) \\ &= G\boldsymbol{p}(\boldsymbol{h})(1-r) - G(K - rF), \end{aligned}$$

where F is the contracted per-unit price of the swap and all other notation is as described in equations 1-3. The final term in the first line represents the hedge payoff in hour h. After rearranging terms to yield the second line in the equation, the only term that contains a quantity not known at the time of contract signing is the first, $G\boldsymbol{p}(\boldsymbol{h})(1-r)$. This term represents the real-time market value of the unhedged portion of the contracted generation, G, in hour h; notably, it vanishes when $r = 1$. That is, when the entire contracted generation quantity is sold forward, the hedge is "perfect" and the buyer bears no risk for the PPA. In this case, the settlement amount in each hour is constant and given by $G(F - K)$. When the buyer is able to sell a swap on the forward market for more than the contracted PPA price, K, this corresponds to a guaranteed profit in each hour of the hedged delivery period.

For the case of a PPA shaped by the P50 hourly profile in each month of the year, the calculation is similar. However, in this case, the hedge ratio must be computed against total monthly contracted generation since the guaranteed shape under the PPA, $G(h)$, varies by hour. Using an ATC swap to hedge some fraction of the total monthly contracted generation will incur marginal "slippage," or misalignment between the hedge payoff and the unhedged PPA settlement amount. More sophisticated hedging programs can be developed to minimize slippage and maximize hedge performance; nonetheless, a significant risk reduction can still be achieved using ATC swaps.

Figure 4

Buyer's August 2019 Cash Flow at Risk (CFaR) for 10 MW Shaped Solar (left) and Wind (right) PPAs as a Function of Hedge Ratio

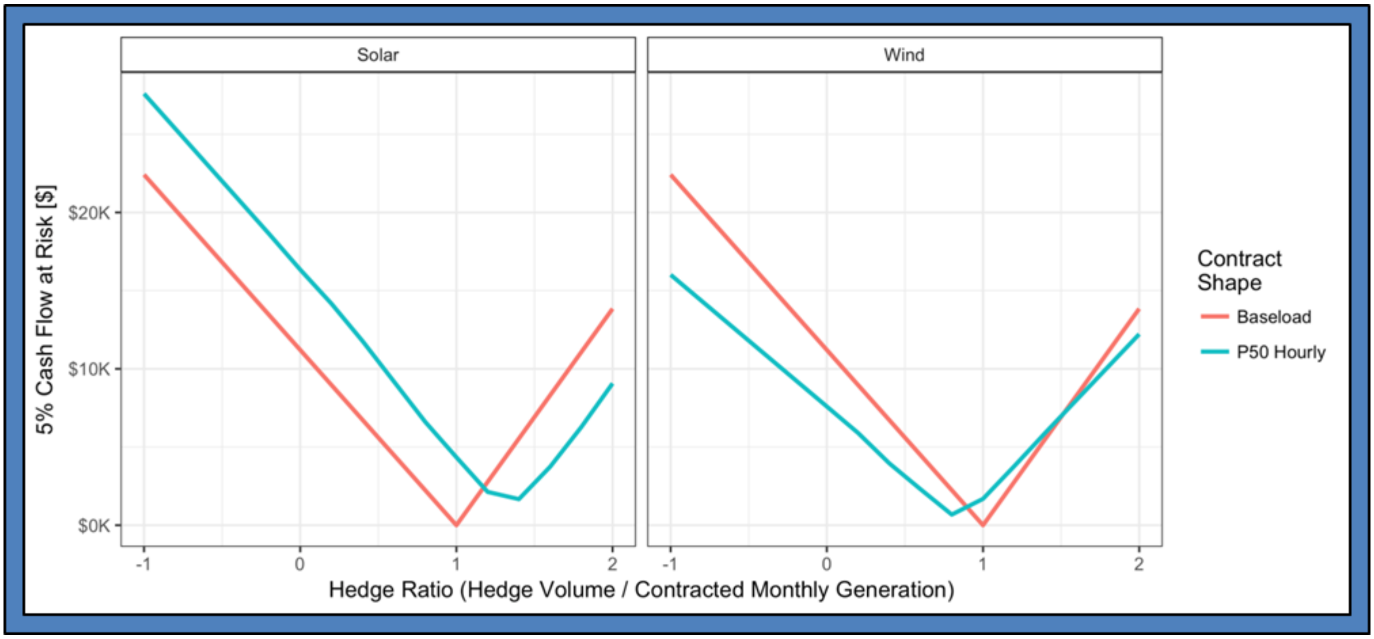

Notes: A hedge ratio of 1.0 corresponds to selling a swap for 100% of the shaped generation volume in the delivery month. Note that the baseload contract is perfectly hedged by selling a swap for full monthly generation volume; in this case, the CFaR is identically zero.

Source of image: cQuant.io ReAssure PPA®.

Figure 4 shows the simulated 5% cash flow at risk (CFaR) for August 2019 for a portfolio containing a long position in a 10 MW PPA and a short swap contracted at various different hedge ratios. The CFaR is computed as the difference between the simulated expected monthly portfolio settlement amount and the simulated P5 settlement amount. Consistent with the theory presented above, the plot shows that the Baseload contracts for both solar and wind (red lines in the figure) are perfectly hedged when 100% of the contracted generation is sold on the forward market. In this case, the CFaR for both the solar and wind PPAs drops to exactly zero and the buyer has a guaranteed profit for each contracted MWh in the amount of the spread between the contract price of the swap, F, and that of the PPA, K.

For the solar and wind PPAs shaped by the P50 hourly generation profile, the CFaR never quite reaches zero for any hedge ratio, though the plots show it can be reduced dramatically from the unhedged case. The optimal hedge ratio for each P50 hourly shaped PPA is the value that corresponds to the lowest point on the blue curve in each of the plots. Interestingly, the optimal hedge quantity is actually greater than one for the solar PPA. This means that the optimal hedge is to sell more than the monthly guaranteed generation under the contract. This seemingly counterintuitive result occurs because the financial product used to hedge the PPA is a short ATC swap, so the hourly hedge volumes are out of alignment with the hourly shaped generation profile of the PPA. This misalignment, in conjunction with the uncertainty in realized generation at the hourly level, creates slippage in the hedge. The result is a distortion of the P50 Hourly CFaR's approach to optimality away from the Baseload case where PPA and hedge volumes align perfectly in each hour.

The interpretation here is that each MWh sold forward in the swap provides slightly less portfolio CFaR reduction than the risk that one MWh of generation under the PPA creates. Thus, more MWh must be sold in the hedge than generated under the PPA in order to minimize risk for the contract. The opposite is true for the wind PPA where the optimal hedge ratio is slightly less than one. This indicates that, up to the optimal hedge quantity, each additional MWh in the hedge provides a slightly greater reduction in risk than the amount that one additional MWh of generation under the PPA would add. In either case, the main takeaway is that the buyer's position in a shaped PPA is relatively straightforward to hedge using standard financial products.[9]

The Seller's Position

Compared to the buyer's position, the seller's risk exposure in a shaped PPA is significantly more complex and contains nuanced interactions between generation and price risk. In addition, these two risk factors cannot be entirely separated, and so must be hedged simultaneously. This becomes clear when looking at the equation for the seller's hourly hedged settlement amount. We again begin by considering settlement amounts for a Baseload shaped PPA, and we assume the seller hedges his native short position with a long ATC swap. Using the notation in equation 4 and adding the swap payoff, the seller's hedged hourly settlement amounts are given by the following equation:

$$\boldsymbol{A_{sell}}(\boldsymbol{h}) = GK + \boldsymbol{g}(\boldsymbol{h})\boldsymbol{p}(\boldsymbol{h}) - G\boldsymbol{p}(\boldsymbol{h}) + rG[\boldsymbol{p}(\boldsymbol{h}) - F]. \#(6)$$

Again, $\boldsymbol{g}(\boldsymbol{h})$ is the actual generation of the facility in hour h, F is the per-unit price of the swap, and we have omitted any explicit dependence on time for the hourly guaranteed quantity, G, since it is the same in all hours under the Baseload contract. The final term in the equation represents the hedge payoff in hour h.

Some minor rearrangement of equation 6 yields a form that is more amenable to discussion:

$$\boldsymbol{A_{sell}}(\boldsymbol{h}) = G\big[K - [rF + \boldsymbol{p}(\boldsymbol{h})(1-r)]\big] + \boldsymbol{g}(\boldsymbol{h})\boldsymbol{p}(\boldsymbol{h}). \#(7)$$

Disassembly of this equation tells us a number of interesting things about the seller's risk profile. First, we note that when the swap price is equal to the PPA price ($F = K$) and the hedge ratio is 100% ($r = 1$), the seller's settlement amount simplifies to $\boldsymbol{A_{sell}}(\boldsymbol{h}) = \boldsymbol{g}(\boldsymbol{h})\boldsymbol{p}(\boldsymbol{h})$, which represents the hourly real-time market value of unit-contingent generation. That is, when the contract price of the swap is equal to the PPA price and the seller hedges 100% of the monthly quantity, his risk exposure is the same as if he were liquidating all production under the PPA directly into the market as-produced. Figure 5 on the next page shows that, in this case, the seller has significantly more risk than in the unhedged case where $r = 0$. This is logical since selling all generation directly into the market as-produced exposes the seller to the full amount of both generation and price risk as well as the magnifying effect of their dynamic interaction. Thus, the PPA itself is a hedge for the seller against his native exposure of renewable generation to real-time electricity prices.

Secondly, the terms inside the outermost pair of square brackets in equation 7 describe the interaction between the PPA price, K, and the value of the shaped energy, G. The total value of shaped energy is determined by the forward market price at the time the hedge is enacted, Fr, and the spot market at delivery, $\boldsymbol{p}(\boldsymbol{h})(1 - r)$. The hedge ratio, r, acts as a risk transfer coefficient, modulating price exposure for the shaped energy between the forward and spot markets. When $r = 1$, the value of the shaped energy in this term is determined entirely by the forward price, F. When $r = 0$, the value is determined entirely by the spot price at delivery, $\boldsymbol{p}(\boldsymbol{h})$, which is the native position without any hedge at all. For values of r between zero and one, value (and risk) is transferred from the spot to the forward market. The practical meaning of this term is that if the forward market moves downward, increasing the hedge ratio by buying additional swaps allows the PPA seller to lock in additional value for the shaped generation amount.

While exploring limiting cases of equation 7 is instructive to build an intuitive understanding of the value and risk in the seller's position, the nonlinear term, $\boldsymbol{g}(\boldsymbol{h})\boldsymbol{p}(\boldsymbol{h})$, makes it impossible to describe the optimal hedge ratio using a closed-form equation. Both the future generation, $\boldsymbol{g}(\boldsymbol{h})$, and the future electricity price, $\boldsymbol{p}(\boldsymbol{h})$, are random variables with distributions that can only be understood empirically. Their product is an even more complex tangle of uncertainty that can only be accessed prior to delivery via numerical simulation. To this end, Figure 5 shows the seller's simulated CFaR as a function of hedge ratio for the same solar and wind PPAs as in Figure 4.

Figure 5
Seller's August 2019 Cash Flow at Risk (CFaR) for 10 MW Shaped Solar (left) and Wind (right) PPAs as a Function of Hedge Ratio

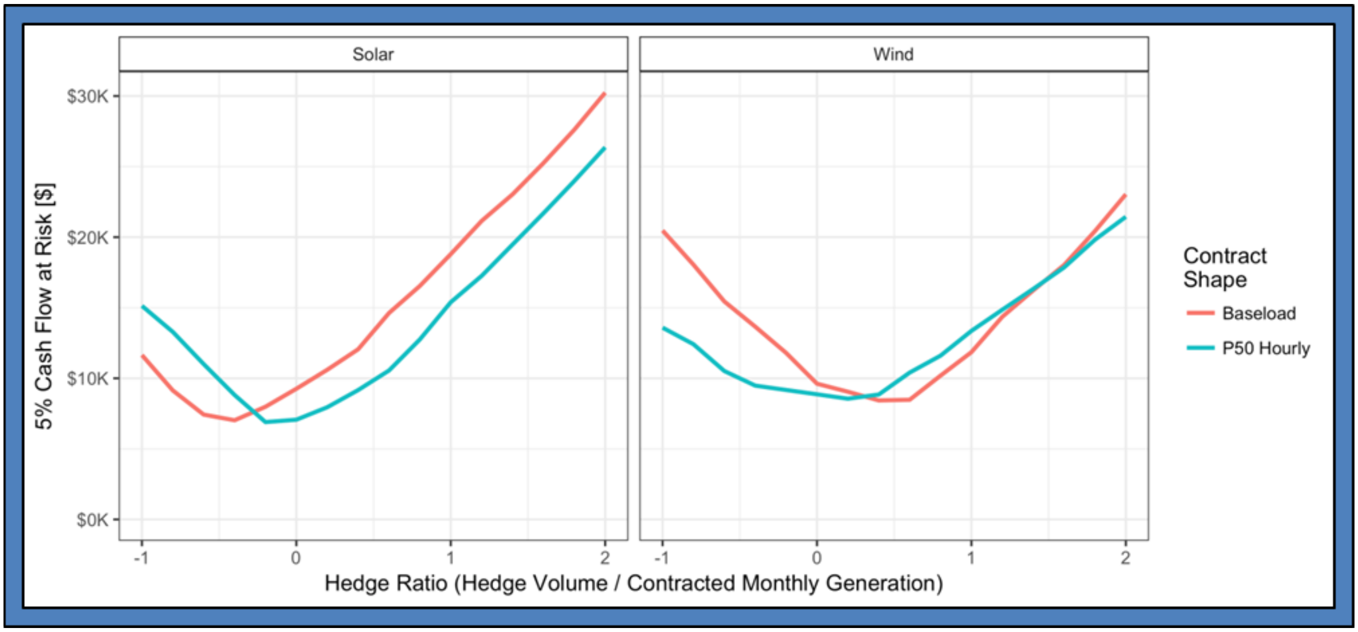

Note: A hedge ratio of 1.0 corresponds to buying a swap for 100% of the shaped generation volume in the delivery month.

Source of image: cQuant.io ReAssure PPA®.

The complicated nature of the seller's risk profile can be seen in the complexity of the curves. Even for the baseload contract shape, the CFaR's approach to optimality (the minimum point on the curve) is far from linear, as it was for the buyer. In all cases, the optimal hedge ratio is far from 100% of the shaped monthly generation; in fact, it is actually negative for solar, indicating a short position should be taken in the swap rather than a long position. We discuss this in more detail below. Finally, it is important to note that the seller's CFaR remains relatively high even when the optimal hedge ratio is used to mitigate the seller's risk for the PPA. Depending on the type of shaping, the best risk reduction the seller can expect using an ATC swap is about a 30% drop in CFaR; in the case of P50 Hourly shaped solar and wind PPAs, the risk reduction is almost nonexistent using this particular financial instrument. This is an indication that more sophisticated hedging schemes and financial products are needed to effectively reduce the seller's risk in these PPAs.

The fact that the optimal hedge ratio for a seller of a baseload solar PPA represents a sale of even more energy on the forward market may appear somewhat counterintuitive initially. The explanation is related to the misalignment of actual solar generation during the delivery month and the guaranteed shape of the PPA. Because of the baseload shape, the seller guarantees the buyer energy during a large number of hours at night when the sun is not shining. During these hours, he must purchase sufficient energy from the spot or forward market to satisfy his contractual obligation to deliver the baseload shape under the PPA.[10] In these nighttime hours, the PPA seller's native position in the spot market flips from long to short; thus, the natural hedge to transfer his spot market exposure to the forward market is to short-sell a swap.

When the seller of the PPA further sells an ATC swap on top of his native short position, he effectively transfers some long exposure to on-peak prices to a corresponding amount of short exposure to off-peak prices. That is, the hedge reduces some of the seller's long spot market exposure in the extremely volatile on-peak hours when the solar facility is likely to be generating above the guaranteed baseload amount. In exchange, the hedge increases the seller's short position in the off-peak hours, requiring him to buy back even more energy from the spot market during these hours to satisfy his obligation to deliver the baseload energy shape. However, because off-peak prices are so much less volatile (uncertain) than on-peak prices, the net result of the hedge is an overall reduction in CFaR for the seller's monthly settlement cash flows. Essentially, the seller is able to gain an overall reduction in cash flow uncertainty by using a short ATC swap to shift a portion of his risk from the on-peak to the off-peak period. This shift in risk comes with a corresponding change in position for the shifted energy during hours when the solar facility's generation is insufficient to satisfy the baseload shape requirement.

Conclusions

The recent increase in adoption of renewable PPAs by a broadening range of market participants has been accompanied by a corresponding increase in the sophistication of these contracts. Some innovative new PPA structures have focused on contractually shaping the delivered energy rather than forcing the buyer to accept both the generation and price risk for the renewable generation. In turn, these new contract structures have opened the door for new strategies to reduce risk for both the buyer and the seller.

In this second article in a two-part series on renewable PPA analytics, valuation, and risk assessment, we discussed how PPA value is built up from interactions between generation and price at the annual, monthly, and hourly level. We also demonstrated how uncertainty in both realized renewable generation and the real-time market price of electricity can have a significant effect on contract value, and how to compute this value through data-driven contract-specific analysis. Building on the discussion of contract shaping, we presented a practical methodology for buyers and sellers to consider how they can hedge their risk(s) in renewable PPAs. We demonstrated that for a particular contract structure at a particular location, we can derive optimal hedge quantities that minimize cash flow at risk for monthly settlement amounts with regard to generation and price risk.

While the present discussion has attempted to provide an introduction to the concepts of shaping and hedging for renewable PPAs, it is far from a comprehensive treatment of the subject. It outlines a set of analytical methodologies that can be leveraged by both buyers and sellers of PPAs to understand value and risk within the contracts and to take action to mitigate risks and prevent loss of value. The methods discussed here can be applied to arbitrary contracted energy shapes and a broad set of hedge instruments. As such, the analysis can be used to generate a detailed understanding of the interplay between the physical energy generated under the PPA and the financial value of that energy within a broader portfolio context and amid complex and volatile energy markets.

Endnotes

1 The first article in the series, entitled "Lifting the Veil on Hidden Risk in Renewable Power Purchase Agreements," provides a more thorough background on renewable PPAs in general and best practices for valuation and risk assessment. The article is on pages 29 to 44 of the Summer 2018 edition of the *GCARD*, which is publicly available and can be downloaded at the following URL: http://www.jpmcc-gcard.com/wp-content/uploads/2018/05/GCARD-Summer-2018.pdf.

2 The terms, contract for differences (CFD), virtual PPA, and synthetic PPA all refer to the same contract structure and may be used interchangeably.

3 Our approach to simulating prices and generation follows the methodology outlined in the first article in this two-part series. Please see "Lifting the Veil on Hidden Risk in Renewable Power Purchase Agreements" in the Summer 2018 edition of the *GCARD*.

4 In practice, if generation is contracted at a particular quantile, it is typically the P99 rather than the P50 (median). The P99 hourly shape provides 99% confidence that the renewable facility will generate at least the shaped amount of energy in each hour. Using the P99 shape gives financiers more confidence that the contracted generation will be produced by the facility and the seller will not have to go to market to make up a deficit, which would create additional risk in the contract. Here we use the P50 because it is a more central statistic and aligns better with our goal of isolating the incremental PPA value of different shaping granularities.

5 Unit-contingent contracts do often include a mechanical availability guarantee in place of a shape guarantee. The availability guarantee removes some or all of the risk to the buyer of mechanical failures or other non-weather-related production deficiencies that may occur throughout the life of the contract. Valuing such an availability guarantee is outside the scope of this article and depends on the probability of incurring an insurable loss in addition to the dynamics of real-time energy generation and market dynamics.

6 A fear of scarcity pricing may even be seen within the forward and futures markets in ERCOT. In the first quarter of 2018, the forward contract for delivery of August 2018 ERCOT North Hub on-peak power traded above $200/MWh in response to announcements of retiring baseload capacity and forecasts of thin summer reserve margins. As of the time this article was written, the average realized real-time price during on-peak hours in the August delivery month was just over $43/MWh. This indicates that the price spike in the forward market may have been largely driven by fear of a possible scarcity pricing event that never materialized.

7 The ATC forward contracts used in the example are assumed to deliver the same quantity in all hours of the delivery month. In practice, on- and off-peak financial products are more actively traded than ATC products. The on- and off-peak periods are defined by trading conventions in the markets where the contracts deliver energy. For example, forward contracts for on-peak power at ERCOT North Hub, as quoted by CME Group, deliver during hours-ending 0700-2200 Monday through Friday, excluding North American Electric Reliability Corporation holidays. Here we have used around-the-clock contracts simply to condense discussion. Please see the following link for additional details on the CME ERCOT North on-peak futures contract and current quoted prices: https://www.cmegroup.com/trading/energy/electricity/ercot-north-zone-mcpe-5-mw-peak-swap-futures.html.

8 In practice, most hedges will incur some degree of basis risk, shape risk, counterparty credit risk, and/or other risks due to imperfect alignment between the native portfolio position and the actively-traded derivative contracts available to hedge it. While these risks too can be mitigated, a thorough discussion of these nuances is outside the scope of the current discussion.

9 The analysis here does not consider margin call risk incurred by exchange-traded futures contract positions, counterparty credit risk incurred by positions in over-the-counter forward contracts, or other risks related to the hedge positions themselves. These risks should be considered within a holistic view of a PPA hedging program.

10 In practice, sellers of a shaped PPA will not typically maintain a short position over the life of the contract. They will either hedge this position in the forward market or will contract for firming and shaping with third-party market participants.

Disclaimer: Although the authors have made every effort to ensure that the information in this article was correct at time of writing, the authors do not assume and hereby disclaim any liability to any party for any loss, damage, or disruption caused by errors or omissions, whether such errors or omissions result from negligence, accident, or any other cause.

This article was prepared by the authors in their personal capacity. The opinions expressed in this article are the authors' own and do not reflect the view of the authors' employers or of the University of Colorado Denver Business School.

Reference

O'Neill, P., 2017, "Fear and Heat in the Texas Power Markets: A Tail-Risk Example and Perspective," *Global Commodities Applied Research Digest*, Editorial Advisory Board Commentaries, Vol. 2, No. 1, Spring, pp. 101-111.

Author Biographies

BROCK MOSOVSKY, Ph.D.
Director of Operations and Analytics, cQuant.io

Dr. Brock Mosovsky is Co-Founder and Director of Operations and Analytics at cQuant.io, a leading software-as-a-service provider of energy-focused advanced quantitative analytics. He has over a decade of analytical modeling experience with a focus on risk management, asset valuation, financial engineering, and model validation for the energy industry. Dr. Mosovsky has worked with the country's largest electric utilities and independent power producers to help improve operational and budget certainty amid volatile financial energy markets and has built and validated financial models for more than 100 gigawatts of nameplate generation capacity. He is also an experienced advisor on valuation and risk assessment of structured transactions and exotic derivatives including asset-specific revenue put options, heat rate call options, and complex power purchase and tolling agreements.

Dr. Mosovsky holds a Ph.D. in Applied Mathematics from the University of Colorado Boulder, a B.S./M.A. in Mathematics from Villanova University, and was awarded a U.S. Fulbright Scholarship for study in the Netherlands.

LANCE TITUS
Managing Director, Uniper Global Commodities

Mr. Lance Titus serves as Managing Director for Uniper Global Commodities. He has over twenty years of commodities trading, structuring and risk management experience. He has held senior leadership roles from Wall Street to the energy industry, working for an investment bank to leading utilities, unregulated entities and merchant energy firms. He has transacted in over $15 billion in trading and originated structured transactions in the energy and environmental commodity markets with a sector focus in electricity, natural gas, renewables, carbon and emissions. Mr. Titus has been a featured panelist at Bloomberg's "The Future of Energy Summit" in New York, and also teaches a course at the University of Colorado Denver on the "Foundations of Commodities."

Mr. Titus holds an M.B.A. from the Daniels College of Business at the University of Denver and a B.S. in Finance and Marketing from Clarion University. He serves on the Advisory Board of cQuant.io and is a member of the J.P. Morgan Center for Commodities' Research Council as well as its Advisory Council. In addition, he serves as a member of the *Global Commodities Applied Research Digest*'s Editorial Advisory Board.