

Market Design in a Zero Marginal Cost Intermittent Renewable Future

by

Frank A. Wolak

Director, Program on Energy and Sustainable Development (PESD)

Holbrook Working Professor of Commodity Price Studies

Department of Economics

Stanford University

Stanford, CA 94305-6072

wolak@zia.stanford.edu

Current Draft: October 15, 2020

The basic features of an efficient short-term wholesale market design do not necessarily need to change to accommodate a significantly larger share of zero marginal cost intermittent renewable energy from wind and solar resources. A large share of controllable zero marginal cost generation does not create any additional market design challenge relative to a market with a large share of controllable positive marginal cost generation. Regardless of the technology, generation unit owners must recover their fixed costs from sales of energy, ancillary services, and long-term resource adequacy products.

A larger variance in the hourly amount of energy produced by intermittent resources is the primary market design challenge associated with a zero marginal cost renewable future. The past ten years in California have demonstrated that as the amount of wind and solar generation capacity increases, the variance in hourly energy produced by these resources increases. This increase in supply uncertainty also increases short-term price volatility, which can finance investments in storage and other technologies that allow consumers to shift their withdrawals of grid-supplied energy away from periods when little wind and solar energy is being produced.

An increased risk of large intermittent energy shortfalls and short-term price volatility implies a greater need for risk management activities. Greater short-term intermittent energy supply risk is likely to require accounting for more transmission and generation operating constraints in the day-ahead and real-time energy markets, as well as purchasing more operating reserves and creating additional ancillary service products. Because controllable generation units are likely to have to start and stop more frequently to make up for unexpected renewable energy shortfalls, there will be a greater need to develop short-term pricing approaches that recover the associated start-up and minimum load costs.

The potential for sustained periods of low intermittent energy production creates both a medium and long-term energy supply risk that requires a new long-term resource adequacy mechanism. The traditional capacity-based approach is unlikely to be the least cost mechanism for ensuring that future demand for energy is met. In a zero marginal cost intermittent future, wind and solar resources must hedge their energy supply risk with controllable generation resources in order to maintain long-term resource adequacy. Cross hedging between these technologies accomplishes two goals. First, it can provide the revenue stream necessary for fixed cost recovery by controllable generation units. Second, it ensures that there is sufficient controllable generation capacity to meet demand under all foreseeable future system states, with a high degree of confidence.

The remainder of this article first describes the key features of an efficient short-term wholesale market design: a multi-settlement locational marginal pricing (LMP) market with an automatic local market power mitigation (LMPM) mechanism, which is the standard market design for all short-term markets in the United States (US). This section concludes with a discussion of modifications of this basic design likely to be necessary to accommodate a larger share of intermittent renewables.

The second half of the paper describes a new long-term resource adequacy mechanism for the efficient short-term market design for an electricity supply industry with a large share of zero marginal cost intermittent renewables. I first explain why a wholesale electricity market requires a long-term resource adequacy mechanism. I then describe a mandated standardized long-term contract approach to long-term resource adequacy that provides strong incentives for intermittent renewable resource owners to hedge their energy supply risk with controllable generation resource owners. I argue that this mechanism ensures long-term resource adequacy at a reasonable cost for final consumers while also allowing the short-term wholesale price volatility that can finance investments in storage and other load-shifting technologies necessary to manage a large share of intermittent renewable resources.

1. Short-Term Market Design

More than twenty-five years of international experience with wholesale electricity market design has identified four crucial features of an efficient short-term market design. First is the extent to which the market mechanism used to set dispatch levels and locational prices is consistent with how the grid and generation units actually operate. Second is a financially binding day-ahead market that prices all transmission and generation unit operating constraints expected to be relevant in real-time. Third is an automatic LMPM mechanism that limits the ability of a supplier to influence the price it receives when it possesses a substantial ability to exercise market power. Fourth are retail market policies that foster active participation of final demand in the wholesale market.

The early US wholesale market designs in the PJM Interconnection, ISO New England, California, and Texas employed simplified versions of the transmission network configuration and generation unit operating constraints. Similar market designs currently exist throughout Europe and the rest of the world. They set a single market-clearing price for an hour or half-hour for an entire control area or for large geographic regions, despite the fact that in real-time there are often generation units with offer prices below this market-clearing price not producing electricity and

units with offer prices above this market-clearing price producing electricity. This outcome occurs because of the location of demand and available generation units within the region, and the configuration of the transmission network prevents some of these low-offer-price units from producing electricity and requires some of the high-offer-price units to supply electricity.

This approach to short-term market design provides incentives for suppliers to take actions to exploit the fact that “in real-time physics wins,” rather than offer their resources into the day-ahead market in a manner that minimizes the cost of meeting demand at all locations in the grid in real-time. Instead, suppliers take actions in the simplified day-ahead market that allow them to profit from knowing they will be needed (or not needed) in real-time because of transmission and generation unit operating constraints.

1.1. Locational Marginal Pricing

Starting with PJM in 1998 and ending with Texas in late 2010, all US wholesale markets adopted a multi-settlement locational marginal pricing (LMP) market design that co-optimizes the procurement of energy and ancillary services and includes an automatic local market power mitigation mechanism built into the market software. This design has a day-ahead financial market which satisfies the locational demands for energy and each ancillary service simultaneously for all 24 hours of the following day. A real-time market then operates using the same network model as the day-ahead market adjusted to real-time system conditions. Deviations from purchases and sales in the day-ahead market are cleared using these real-time prices. Both of these markets price all relevant transmission network and other relevant operating constraints on generation units. As I discuss below, this market design can foster active participation of final demand in the wholesale market.

Only generation unit output levels that are physically feasible will be accepted in both the day-ahead and real-time markets. Prices for the same hour vary depending on whether the location is in a generation-deficient or generation-rich region of the transmission network. The locational marginal price or nodal price at a given location is the increase in the minimized value of the “as-offered costs” of serving the locational demands for energy and all ancillary services as a result of a one-unit increase in the amount of energy withdrawn at that location in the transmission network. The price of each ancillary service is equal to the increase in the optimized value of the objective function as a result of a one-unit increase in the demand for that ancillary service.

The recent experience of many European countries with significant wind and solar resources indicates that the cost of making the final schedules that emerge from their zonal markets physically feasible is likely to get even larger as the amount of intermittent renewable generation capacity increases. According to the European Network of Transmission System Operators for Electricity, in 2017 these costs were over 1 billion euros in Germany, more than 400 million euros in Great Britain, over 80 million euros in Spain, and approximately 50 million euros in Italy.

1.2. Multi-Settlement LMP Market

A multi-settlement LMP market has at least a day-ahead forward market and a real-time market, each of which employs the same market-clearing mechanism. The day-ahead market typically allows generation unit owners to submit three-part offers to supply energy: start-up costs, minimum load costs, and an energy offer curve. These are used to compute hourly generation schedules and ancillary service quantities and LMPs for energy and ancillary services for all 24 hours of the following day. A generation unit will not be accepted to supply energy in the day-ahead market unless the combination of its offered start-up costs, minimum load costs and energy production costs are part of the least as-offered-cost solution to serving the hourly locational demands for all 24 hours of the following day.

The energy schedules that emerge from the day-ahead market do not require a generation unit to produce the energy sold or a load to consume the energy purchased in the day-ahead market at a given location. Any production shortfall relative to a day-ahead generation schedule must be purchased from the real-time market at that location. Any production greater than a generation unit's day-ahead schedule is sold at the real-time price at that location. Any additional consumption beyond a load's day-ahead energy schedule is paid for at the real-time price at that location, and the surplus of a day-ahead schedule relative to actual consumption is sold at the real-time price at that location.

1.3. Mitigating Local Market Power

The configuration of the transmission network, the level and location of demand, as well as the level of output of other generation units can create system conditions in which almost any generation unit or group of generation units has a significant ability to exercise unilateral market power. The constrained-on generation problem is an example of this phenomenon. The unit's owner knows that it must be accepted to supply energy regardless of its offer price. Without an LMPM mechanism, there may be no limit to the offer price the unit's owner could submit and be

accepted to supply energy. During the first summer of the California market, when there was no formal LMPM mechanism, suppliers submitted extremely high offers for energy and ancillary services when these system conditions arose.

This logic is why market power mitigation mechanisms typically used in Europe and other industrialized regions and initially employed in the US, that designate in advance the offers of certain generation units for mitigation for an entire year, miss many instances of the exercise of substantial unilateral market power.

An automated LMPM mechanism built into the market software that relies on actual system conditions to determine whether any supplier has a substantial ability and incentive to exercise unilateral market power is likely to be significantly more effective. This regulator-approved administrative procedure determines: (1) when a supplier has an ability to exercise local market power worthy of mitigation, (2) the value of the supplier's mitigated offer price, and (3) the price the mitigated supplier is paid. It is increasingly clear to regulators around the world, particularly those that operate markets with a finite amount of transmission capacity and significant intermittent renewable generation capacity, that an automatic LMPM mechanism is a necessary feature of any short-term market design.

Because these LMPM mechanisms are built into the market software of all US markets and automatically mitigate the offers of suppliers deemed to have substantial ability to exercise unilateral market power, they are effective at preventing the exercise of significant local market power with little disruption to the operation of the short-term market.

1.4. Benefits of a Multi-Settlement LMP Market

A multi-settlement LMP market design can facilitate the active participation of final consumers in the wholesale market and reduce both the input fuel and total variable cost of producing the same amount of thermal energy relative to a multi-settlement zonal market design. The presence of an automatic LMPM mechanism and make-whole payments that guarantee start-up, minimum load, and energy cost recovery for the day for all generation units committed to operate in the day-ahead market reduces the incentive for suppliers to exercise unilateral market power. An expected profit-maximizing supplier with no ability to exercise unilateral market power will submit an offer price equal to its marginal cost because make-whole payments ensure recovery of their start-up, minimum load, and energy costs.

Because day-ahead purchases are firm financial commitments, a retailer can sell energy

purchased in the day-ahead market at the real-time price by consuming less than its day-ahead energy schedule. This eliminates the need for the regulator to set an administrative baseline relative to which a retailer sells demand reductions. The day-ahead market also allows retailers and large consumers to submit price-sensitive bid curves into the day-ahead market to reduce the market-clearing price and the quantity of energy they purchase in the day-ahead market.

1.5. Modifications for Large Scale Intermittent Renewables Deployment

A multi-settlement LMP market design is capable of managing a generation mix with a significant share of intermittent renewables. However, some modifications are likely to be needed as the share of intermittent renewable resources increases. Additional operating constraints will need to be incorporated into the day-ahead and real-time market models for reliable system operation with an increased quantity of intermittent renewables.

There is also likely to be a need to introduce additional ancillary services to accommodate a larger share of intermittent renewable energy. For example, California introduced a fast-ramping ancillary service product that compensates controllable generation units for not supplying energy during certain hours of the day in order to have sufficient unloaded capacity to meet the rapid increase in net demand (the difference between system demand and renewable generation) in the early evening when the state's solar resources stop producing.

Because controllable resources are likely to have to start and stop more frequently as the share of intermittent resources increases, implementations of convex hull pricing and other market-clearing mechanisms that limit the magnitude of make-whole payments will need to be developed.

2. Resource Adequacy with Significant Intermittent Renewables

Why do wholesale electricity markets require a regulatory mandate to ensure long-term resource adequacy? Electricity is essential to modern life, but so are many other goods and services. Consumers want cars, but there is no regulatory mandate that ensures enough automobile assembly plants to produce these cars. They want point-to-point air travel, but there is no regulatory mandate to ensure enough airplanes to accomplish this. Many goods are produced using high fixed cost, low marginal cost technologies, similar to electricity supply. Nevertheless, these firms recover their cost of production, including a return on the capital invested, by selling their output at a market-determined price.

So, what is different about electricity that requires a long-term resource adequacy mechanism? The regulatory history of the electricity supply industry and the legacy technology for metering electricity consumption results in what I call a *reliability externality*.

2.1. The Reliability Externality

Different from the case of wholesale electricity, in the market for automobiles and air travel there is no regulatory prohibition on the short-term price rising to the level necessary to clear the market. Airlines adjust the prices for seats on a flight over time in an attempt to ensure that the number of customers traveling on that flight equals the number of seats flying. This ability to use price to allocate the available seats is also what allows the airline to recover its total production costs.

Using the short-term price to manage the real-time supply and demand balance in a wholesale electricity market is limited by a finite upper bound on a supplier's offer price and/or a price cap that limits the maximum market-clearing price. Although offer caps and price caps can limit the ability of suppliers to exercise unilateral market power in the short-term energy market, they also reduce the revenues suppliers can receive during scarcity conditions. This is often referred to as the *missing money* problem for generation unit owners. However, this missing money problem is only a symptom of the existence of the “reliability externality.”

This externality exists because offer caps limit the cost to electricity retailers of failing to hedge their purchases from the short-term market. Specifically, if the retailer or large consumer knows the price cap on the short-term market is \$250/MWh, then it is unlikely to be willing to pay more than that for electricity in any earlier forward market. This creates the possibility that real-time system conditions can occur where the amount of electricity demanded at or below the offer cap is less than the amount suppliers are willing to offer at or below the offer cap. This outcome implies that the system operator must be forced to either abandon the market mechanism or curtail load until the available supply offered at or below the offer cap equals the reduced level of demand, as occurred a number of times in California between January 2001 and April 2001, and most recently on August 14 and 15, 2020.

Because random curtailments of supply—also known as rolling blackouts—are used to make demand equal to the available supply at or below the offer cap under these system conditions, this mechanism creates a “reliability externality” because no retailer bears the full cost of failing to procure adequate amounts of energy in advance of delivery. A retailer that has purchased

sufficient supply in the forward market to meet its actual demand is equally likely to be randomly curtailed as another retailer of the same size that has not procured adequate energy in the forward market. For this reason, all retailers have an incentive to under-procure their expected energy needs in the forward market.

The lower the offer cap, the greater is the likelihood that the retailer will delay their electricity purchases to the short-term market. Delaying more purchases to the short-term market increases the likelihood of insufficient supply in the short-term market at or below the offer cap. Because retailers do not bear the full cost of failing to procure sufficient energy in the forward market to meet their future demand, there is a missing market for long-term contracts for long enough delivery horizons into the future to allow new generation units to be financed and constructed to serve demand under all future conditions in the short-term market. Therefore, a regulator-mandated long-term resource adequacy mechanism is necessary to replace this missing market.

Some form of regulatory intervention is necessary to internalize the resulting reliability externality, unless the regulator is willing to eliminate or substantially increase the offer cap so that the short-term price can be used to equate available supply to demand under all possible future system conditions. This approach is taken by the Electricity Reliability Council of Texas (ERCOT), which has a \$9,000/MWh offer cap, and National Electricity Market in Australia, which has a 15,000 Australia Dollars per MWh offer cap. If customers do not have interval meters that can record their consumption on an hourly basis, then they have a very limited ability to benefit from shifting their consumption away from high-priced hours. All that can be recorded for these customers is their total consumption between two successive meter readings so they can only be billed based on an average wholesale price during the billing cycle. Therefore, raising or having no offer cap on the short-term market would not be advisable in a region where few customers have interval meters. Even in regions with interval meters, there would be substantial political backlash from charging hourly wholesale prices that cause real-time demand to equal available supply under all possible future system conditions.

Currently, the most popular approach to addressing this reliability externality is a capacity payment mechanism that assigns a firm capacity value to each generation unit based on the amount of energy it can provide under stressed system conditions. Sufficient firm capacity procurement obligations are then assigned to retailers to ensure that annual system demand peaks can be met.

Capacity-based approaches to long-term resource adequacy rely on the credibility of the firm capacity measures assigned to generation units. This is a relatively straightforward process for thermal units. The nameplate capacity of the generation unit times its annual availability factor is a reasonable estimate of the amount of energy the unit can provide under stressed system conditions. For the case of hydroelectric facilities, this process is less straightforward. The typical approach uses percentiles of the distribution of past hydrological conditions for that generation unit to determine its firm capacity value.

Assigning a firm capacity value to a wind or solar generation unit is extremely challenging for several reasons. First, these units only produce when the underlying resource is available. If stressed system conditions occur when the sun is not shining or the wind is not blowing, these units should be assigned little, if any, firm capacity value. Second, because there is a high degree of contemporaneous correlation between the energy produced by solar and wind facilities within the same region, the usual approach to determining the firm capacity of a wind or solar unit assigns a smaller value to that unit as the total MWs of wind or solar capacity in the region increases. These facts imply that a capacity-based long-term resource adequacy mechanism is poorly suited to a zero marginal cost intermittent renewable feature.

2.2. Supplier Incentives with Fixed-Price Forward Contract Obligations for Energy

The standardized fixed-price forward contract (SFPFC) approach to long-term resource adequacy recognizes that a supplier with the ability to serve demand at a reasonable price may not do so if it has the ability to exercise unilateral market power. A supplier with the ability to exercise unilateral market power with a fixed-price forward contract obligation finds it expected profit maximizing to minimize the cost of supplying this forward contract quantity of energy. The SFPFC long-term resource adequacy mechanism takes advantage of this incentive by requiring retailers to hold hourly fixed-price forward contract obligations for energy that sum to the hourly value of system demand. This implies that all suppliers find it expected profit maximizing to minimize the cost of meeting their hourly fixed-price forward contract obligations, the sum of which equals the hourly system demand for all hours of the year.

To understand the logic behind the SFPFC mechanism, consider the example of a supplier that owns 150 MWs generation capacity that has sold 100 MWh in a fixed-forward contract at a price of \$25/MWh for a certain hour of the day. This supplier has two options for fulfilling this forward contract: (1) produce the 100 MWh energy from its own units at their marginal cost of

\$20/MWh or (2) buy this energy from the short-term market at the prevailing market-clearing price. The supplier will receive \$2,500 from the buyer of the contract for the 100 MWh sold, regardless of how it is supplied. This means that the supplier maximizes the profits it earns from this fixed-price forward contract sale by minimizing the cost of supplying the 100 MWh of energy.

To ensure that the least-cost “make versus buy” decision for this 100 MWh is made, the supplier should offer 100 MWh in the short-term market at its marginal cost of \$20/MWh. This offer price for 100 MWh ensures that if it is cheaper to produce the energy from its generation units—the market price is at or above \$20/MWh—the supplier’s offer to produce the energy will be accepted in the short-term market. If it is cheaper to purchase the energy from the short-term market—the market price is below \$20/MWh—the supplier’s offer will not be accepted and the supplier will purchase the 100 MWh from the short-term market at a price below \$20/MWh.

This example demonstrates that the SFPFC approach to long-term resource adequacy makes it expected profit-maximizing for each seller to minimize the cost supplying the quantity of energy sold in this forward contract each hour of the delivery period. By the logic of the above example, each supplier will find it in its unilateral interest to submit an offer price into the short-term market equal to its marginal cost for its hourly SFPFC quantity of energy, in order to make the efficient “make versus buy” decision for fulfilling this obligation.

In addition, because all suppliers know that the sum of the values of the hourly SFPFC obligations for all suppliers is equal the system demand, each firm knows that its competitors have substantial fixed-price forward contract obligations for that hour. This implies that all suppliers know that they have limited opportunities to raise the price they receive for short-term market sales beyond their hourly SFPFC quantity. For the above example, the supplier that owns 150 MWs of generation capacity has a strong incentive to submit an offer price close to its marginal cost to supply any energy beyond the 100 MWh of SFPFC energy it is capable of producing. Therefore, attempts by any supplier to raise prices in the short-term market by withholding output beyond their SFPFC quantity are likely to be unsuccessful because of the aggressiveness of the offers into the short-term market by its competitors with hourly SFPFC obligations.

2.3. SFPFC Approach to Resource Adequacy

This long-term resource adequacy mechanism requires all electricity retailers to hold SFPFCs for energy for fractions of realized system demand at various horizons to delivery. For example, retailers in total must hold SFPFCs that cover 100 percent of realized system demand,

95 percent of realized system demand one year in advance of delivery, 90 percent two-years in advance of delivery, 87 percent three years in advance of delivery, and 85 percent four years in advance of delivery. The fractions of system demand and number of years in advance that the SFPFCs must be purchased are parameters set by the regulator to ensure long-term resource adequacy. In the case of a multi-settlement LMP market, the SFPFCs would clear against the quantity-weighted average of the hourly locational prices at all load withdrawal nodes.

SFPFCs are shaped to the hourly system demand within the delivery period of the contract. Figure 1 contains a sample pattern of system demand for a four-hour delivery horizon. The total demand for the four hours is 1000 MWh, and the four hourly demands are 100 MWh, 200 MWh, 400 MWh and 300 MWh. Therefore, a supplier that sells 300 MWh of SFPFC energy has the hourly system demand-shaped forward contract obligations of 30 MWh in hour 1, 60 MWh in hour 2, 120 MWh in hour 3 and 90 MWh in hour 4 for Firm 1 in Figure 2. The hourly forward contract obligations for Firm 2 that sold 200 MWh SFPFC energy and Firm 3 that sold 500 MWh of SFPFC energy are also shown in Figure 2. These SFPFC obligations are also allocated across the four hours according to the same four hourly shares of total system demand. This ensures that the sum of the hourly values of the forward contract obligations for the three suppliers is equal to the hourly value of system demand. Taking the example of hour 3, Firm 1's obligation is 120 MWh, Firm 2's is 80 MWh and Firm 3's is 200 MWh. These three values sum to 400 MWh, which is equal to the value of system demand in hour 3 shown in Figure 1.

These standardized fixed-price forward contracts are allocated to retailers based on their share of system demand during the month. Suppose that the four retailers in Figure 3 consume $1/10$, $2/10$, $3/10$, and $4/10$, respectively, of the total energy consumed during the month. This means that Retailer 1 is allocated 100 MWh of the 1000 MWh SFPFC obligations for the four hours, Retailer 2 is allocated 200 MWh, Retailer 3 is allocated 300 MWh, and Retailer 4 is allocated 400 MWh. The obligations of each retailer are then allocated to the individual hours using the same hourly system demand shares used to allocate the SFPFC energy sales of suppliers to the four hours. This allocation process implies Retailer 1 holds 10 MWh in hour 1, 20 MWh in hours 2, 40 MWh in hour 3 and 30 MWh in hour 4. Repeating this same allocation process for the other three retailers yields the remaining three hourly allocations shown in Figure 3. Similar to the case of the suppliers, the sum of allocations across the four retailers for each hour equals the total hourly system demand. For period 3, Retailer 1's holding is 40 MWh, Retailer 2's is 80 MWh, Retailer 3's is 120 MWh, and Retailer 4's is 160 MWh. The sum of these four magnitudes

is equal to 400 MWh, which is the system demand in hour 3.

2.4. Mechanics of Standardized Forward Contract Procurement Process

The SFPFCs would be purchased through auctions several years in advance of delivery in order to allow new entrants to compete to supply this energy. Because the aggregate hourly values of these SFPFC obligations are allocated to retailers based on their actual share of system demand during the month, this mechanism can easily accommodate retail competition. If one retailer loses load and another gains it during the month, the share of the aggregate hourly value of SFPFCs allocated to the first retailer falls and the share allocated to the second retailer rises.

The wholesale market operator would run the auctions with oversight by the regulator. One advantage of the design of the SFPFC products is that a simple auction mechanism can be used to purchase each annual product. A multi-round auction could be run where suppliers submit the total amount of annual SFPFC energy they would like to sell for a given delivery period at the price for the current round. Each round of the auction the price would decrease until the amount suppliers are willing to sell at that price is less than or equal to the aggregate amount of SFPFC energy demanded.

The wholesale market operator would also run a clearinghouse to manage the counterparty risk associated with these contracts. All US wholesale market operators currently do this for all participants in their energy and ancillary services markets. In several US markets, the market operator also provides counterparty risk management services for long-term financial transmission rights, which is not significantly different from performing this function for SFPFCs.

SFPFCs auctions would be run on an annual basis for deliveries starting two, three, and four years in the future. In steady state, auctions for incremental amounts of each annual contract would also be needed so that the aggregate share of demand covered by each annual SFPFC could increase over time. The eventual 100 percent coverage of demand occurs through a final true-up auction that takes place after the realized values for hourly demand for the delivery period are known.

Consider the following two examples of how the true-up auction would work. Assume for simplicity, the monthly load shares of the four retailers remain unchanged. Suppose that the initial 1000 MWh SFPFC in the above example sold at \$50/MWh. However, suppose that actual demand turned out to be 10 percent higher in every period as shown Figure 4 and the additional 100 MWh purchased in the true-up auction sold at \$80/MWh. If each firm sold 10 percent more SFPFC

energy in the true-up auction this would yield the hourly obligations for each supplier shown in Figure 5. The hourly obligations for the four retailers are shown in Figure 6. These would clear against the average cost of purchases from the original auction and true-up auction of \$52.73. If the realized hourly demands are ten percent lower as shown in Figure 7, the true-up auction would buy back 100 MWh of SFPFC energy. If all suppliers bought back 10 percent of their initial sales at \$20/MWh, the resulting hourly obligations would be those shown in Figure 8. The 10 percent smaller hourly obligations of the four retailers are shown in Figure 9 and these would clear against the average cost of the initial auction purchase less the revenues from the true-up auction sales for the required 900 MWh of obligations of \$53.33.

As shown in Figures 6 and 9, each purchase or sale of the same annual SFPFC product is allocated to retailers according to their load shares during the delivery month. If three different size purchases are made for the same annual SFPFC product at different prices, then each retailer is allocated its load share for the month of these three purchases. This ensures a level playing field for retailers with respect to their long-term resource adequacy obligation. All retailers face the same average price for the long-term resource adequacy obligation associated with their realized demand for the month.

The advance purchase fractions of the final demand are the regulator's security blanket to ensure that system demands can be met for all hours of the year for all possible future system conditions. If the regulator is worried that not enough resources will be available in time to satisfy this requirement, it can increase the share of final demand that it purchases in each annual SFPFC auction. As shown above, if too much SFPFC energy is purchased in an annual auction, it can be sold back to generation unit owners in a later auction or the final true-up auction.

Cross hedging between controllable generation units and intermittent renewable resources under this mechanism is enforced by tying the amount of SFPFC energy a generation unit owner can sell on an annual basis to the value of their firm energy. The system operator would assign firm energy values for each generation unit using a mechanism similar to what is currently used to compute firm capacity values. Multiplying a unit's MWs of firm capacity by the number of hours in the year would yield the unit's firm energy value, which is the upper bound on the amount of SFPFC energy the unit owner could sell in all auctions for an annual compliance period. Because the firm capacity of a generation unit is defined as the amount of energy it can produce under stressed system conditions, this limitation on annual sales of firm energy implies that intermittent wind and solar resources would sell much less SFPFC energy than the total MWhs they expect to

produce in a typical year, and controllable generation unit owners would sell significantly more SFPFC energy than the total MWhs they expect to produce in an typical year.

In most years, a controllable resource owner would be producing energy in a small number of hours of the year, but earning the difference between the price at which they sold the energy in the SFPFC auction and the hourly short-term market price times the hourly value of its SFPFC energy obligation for all the hours that it does not produce energy. Intermittent renewables owners would typically produce more than their SFPFC obligation in energy and sell the additional energy at the short-term price. In years with low renewable output near their SFPRC obligations, controllable resource owners would produce close to the hourly value of their SFPFC energy obligation, thus making average short-term prices significantly higher. However, aggregate retail demand would be shielded from these high short-term prices because of their SFPC holdings.

2.4. Advantages of SFPFC Approach to Long-Term Resource Adequacy

This mechanism has a number of advantages relative to a capacity-based approach. There is no regulator-mandated aggregate capacity requirement. Generation unit owners are allowed to decide both the total MWs and the mix of technologies to meet their SFPFC energy obligations. There is also no prohibition on generation unit owners or retailers engaging in other hedging arrangements outside of this mechanism. Specifically, a retailer could enter into a bilateral contract for energy with a generation unit owner or other retailer to manage the short-term price and quantity risk associated with the difference between their actual hourly load shape and the hourly values of their retail load obligation. This mechanism provides a nudge to market participants to develop a liquid market for these bilateral contract arrangements at horizons to delivery similar to the SFPFC products. Instead of starting from the baseline of no fixed-price forward contract coverage of system demand by retailers, this mechanism starts with 100 percent coverage of system demand, which retailers can unwind at their own risk.

For the regulated retail customers, the purchase prices of SFPFCs can be used to set the wholesale price implicit in the regulated retail price over the time horizon that the forward contract clears. This would provide retailers with a strong incentive to reduce their average wholesale energy procurement costs below this price through bilateral hedging arrangements, storage investments, or demand response efforts.

There are several reasons why this mechanism should be a more cost-effective approach to long-term resource adequacy than a capacity-based mechanism in a zero marginal cost intermittent

future. First, the sale of SFPFC energy starting delivery two or more years in the future provides a revenue stream that will significantly increase investor confidence in recovering the cost of any investment in new generation capacity.

Second, because retailers are protected from high short-term prices by total hourly SFPFC holdings equal to system demand, the offer cap on the short-term market can be raised in order to increase the incentive for all suppliers to produce as much energy as possible during stressed system conditions. Third, the possibility of higher short-term price spikes can finance investments in storage and load-shifting technologies and encourage active participation of final demand in the wholesale market, further enhancing system reliability in a market with significant intermittent renewable resources.

If SFPFC energy is sold for delivery in four years based on a proposed generation unit, the regulator should require construction of the new unit to begin within a pre-specified number of months after the signing date of the contract or require posting of a substantially larger amount of collateral in the clearinghouse with the market operator. Otherwise, the amount of SFPFC energy that this proposed unit sold would be automatically liquidated in a subsequent SFPFC auction and a financial penalty would be imposed on the developer. Other completion milestones would have to be met at future dates to ensure the unit is able to provide the amount of firm energy that it committed to provide in the SFPFC contract sold. If any of these milestones were not met, the contract would be liquidated.

3. Final Comments

There is no perfect wholesale market design. There are only better wholesale market designs, and what constitutes a better design depends on many factors specific to the region. The long-term resource adequacy mechanism should be coordinated with the short-term market design. Although there is general agreement on the key features of a best-practice short-term market design, many details must be adjusted to reflect local conditions. For this reason, wholesale market design is a process of continuous learning, adaption, and hopefully, improvement. The standardized energy contracting approach to long-term resource adequacy described in this paper is an example of this process. While it has a number of features that are likely make it significantly better suited to a zero-marginal-cost intermittent-renewables-electricity-supply industry, there are many details of this basic mechanism that should be adapted to reflect local conditions.

Further Reading

Graf, C., Quaglia, F. and Wolak, F.A. (2020) “Simplified Electricity Market Models with Significant Intermittent Renewable Capacity: Evidence from Italy (January 2020), available at <http://www.stanford.edu/~wolak>.

Gribik, P. R., Hogan, W. W., & Pope, S. L. (2007). Market-clearing electricity prices and energy uplift. *Cambridge, MA*.

Hogan, W. W. (2002). Electricity market restructuring: reforms of reforms. *Journal of Regulatory Economics*, 21(1), 103-132.

Schweppe, F. C., Caramanis, M. C., Tabors, R. D., & Bohn, R. E. (2013). *Spot pricing of electricity*. Springer Science & Business Media.

Wolak, F. A. (2020) Wholesale Electricity Market Design, in *Handbook on the Economics of Electricity Markets*, J.M. Glanchant, P.L. Joskow, and M. Pollitt (editors), Cambridge University Press.

Biography

Frank A. Wolak is with Stanford University

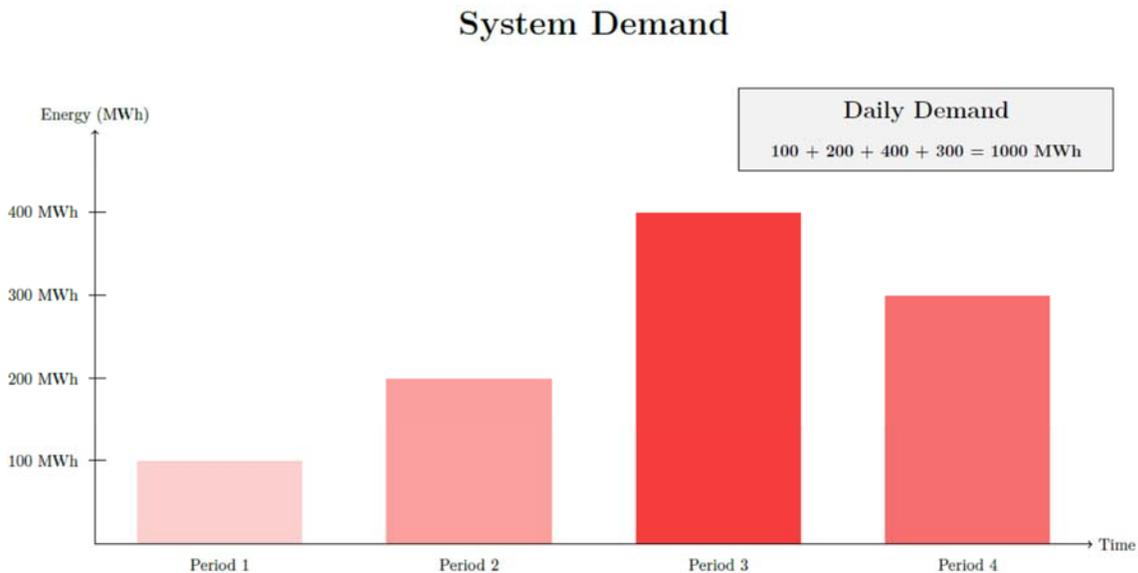


Figure 1: Hourly System Demands

Three Firms:
Firm 1 sells 300 MWh
Firm 2 sells 200 MWh
Firm 3 sells 500 MWh
Total Amount Sold by Three Firms = 1000 MWh

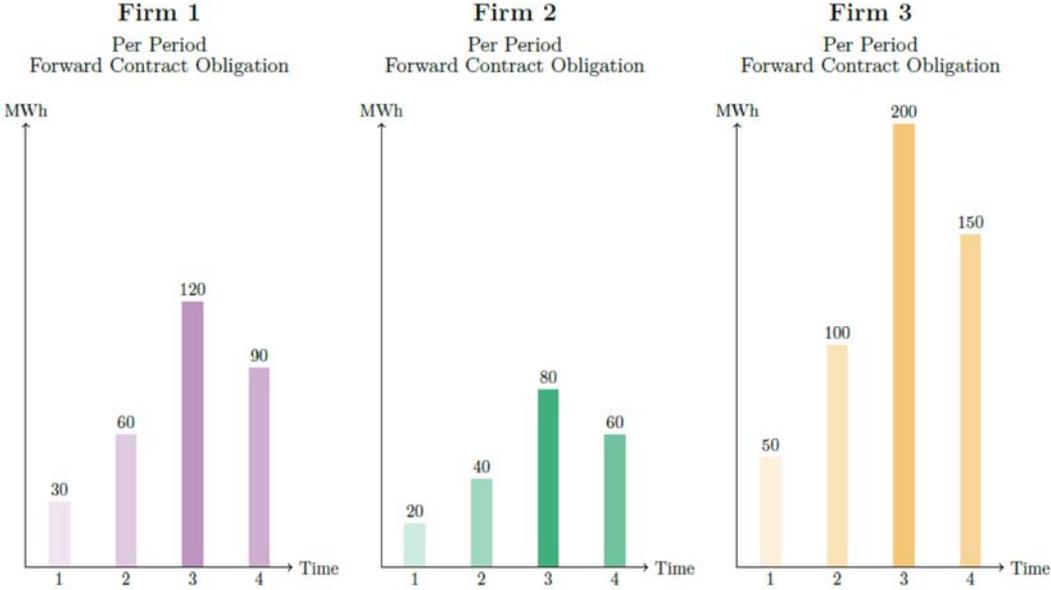


Figure 2: Hourly Forward Contract Quantities for Three Suppliers

Four Retailers:
 Retailer 1 holds 100 MWh
 Retailer 2 holds 200 MWh
 Retailer 3 holds 300 MWh
 Retailer 4 holds 400 MWh
 Total Amount Held by Four Retailers = 1000 MWh



Figure 3: Hourly Forward Contract Quantities for Four Retailers

System Demand

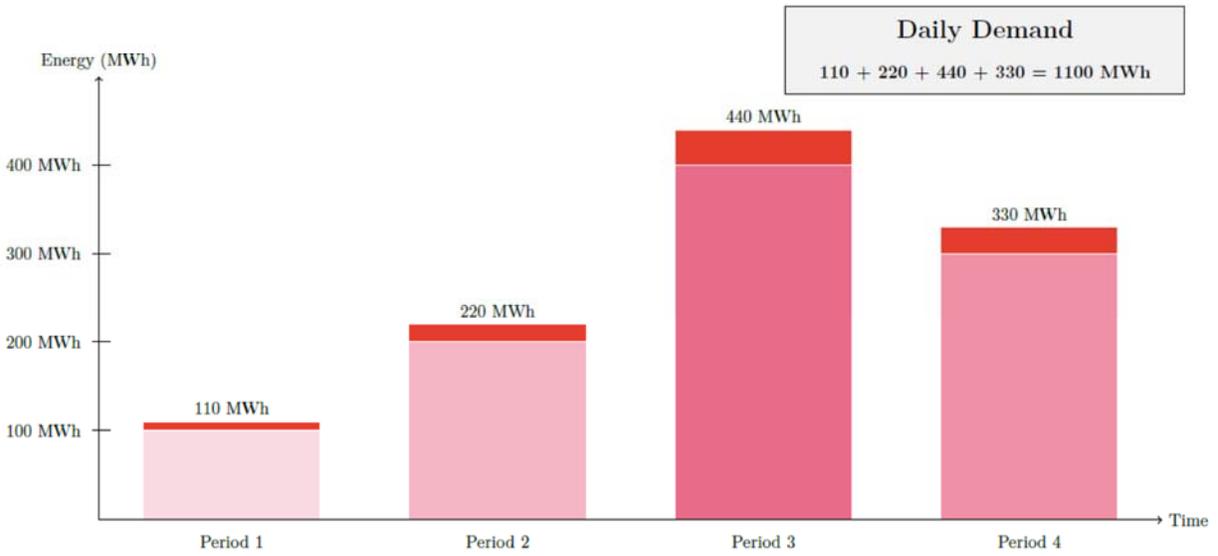


Figure 4: Hourly System Demands (10 Percent Higher)

Three Firms:
 Firm 1 sells 330 MWh
 Firm 2 sells 220 MWh
 Firm 3 sells 550 MWh
 Total Amount Sold by Three Firms = 1100 MWh



Figure 5: Hourly Forward Contract Quantities for Three Suppliers (10 Percent Higher)

Four Retailers:
 Retailer 1 holds 110 MWh
 Retailer 2 holds 220 MWh
 Retailer 3 holds 330 MWh
 Retailer 4 holds 440 MWh
 Total Amount Held by Four Retailers = 1100 MWh



Figure 6: Hourly Forward Contract Quantities for Four Retailers (10 Percent Higher)

System Demand

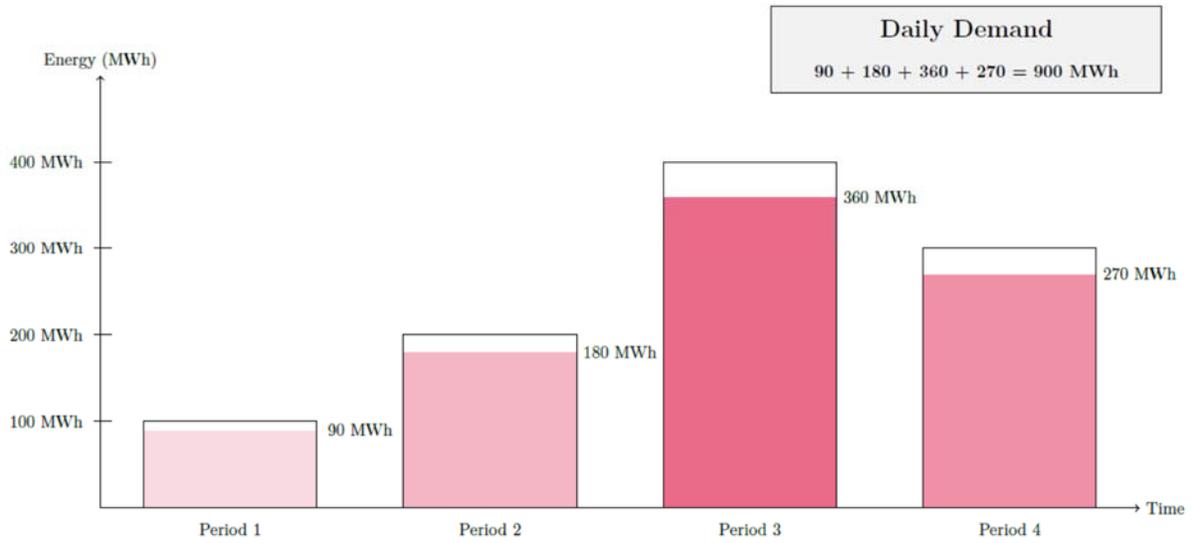


Figure 7: Hourly System Demands (10 Percent Lower)

Three Firms:
 Firm 1 sells 270 MWh
 Firm 2 sells 180 MWh
 Firm 3 sells 450 MWh
 Total Amount Sold by Three Firms = 900 MWh

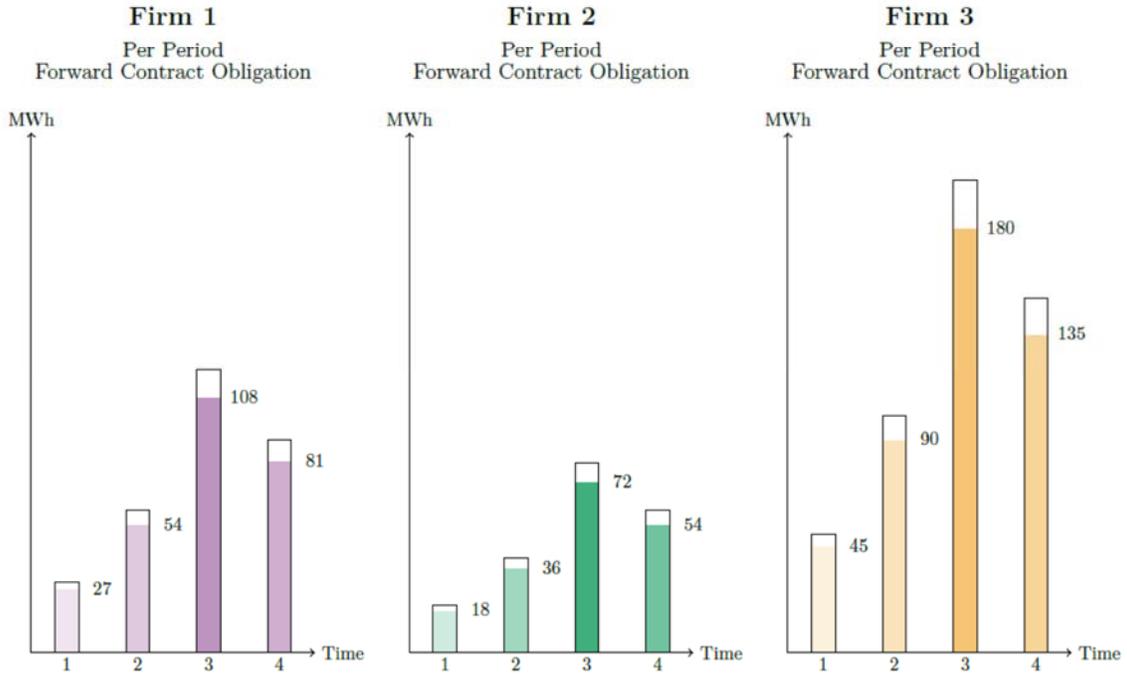


Figure 8: Hourly Forward Contract Quantities for Three Suppliers (10 Percent Lower)

Four Retailers:
 Retailer 1 holds 90 MWh
 Retailer 2 holds 180 MWh
 Retailer 3 holds 270 MWh
 Retailer 4 holds 360 MWh
 Total Amount Held by Four Retailers = 900 MWh

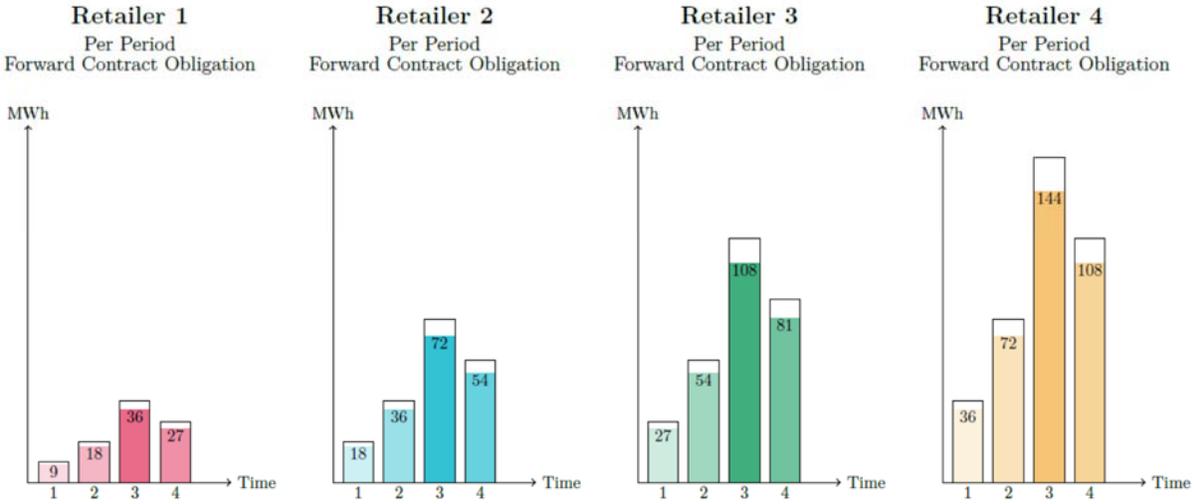


Figure 9: Hourly Forward Contract Quantities for Four Retailers (10 Percent Lower)