

Electricity Market Design: The Good, the Bad, and the Ugly

Peter Cramton
University of Maryland
peter@cramton.umd.edu

Abstract

This paper examines principles of market design as applied to electricity markets. I illustrate the principles with examples of both good and bad designs. I discuss one of the main design challenges—dealing with market power. I then discuss FERC's choice of a standard market design.

1. Introduction

In recent years, the use of complex auctions to allocate scarce resources has grown tremendously. The growth stems in part from technological advances, which greatly reduce the cost of conducting such auctions, and in part from the deregulation of several infrastructure industries, which puts a premium on using market-based methods to allocate resources. This paper examines the recent experience with high-stake auctions in deregulated industries, especially energy and telecommunications. Lessons in market design are drawn from this experience.

Good market design is crucial to the success of these markets. Good market design identifies the critical issues, and then addresses them as simply as possible, but not more simply. For example, the simultaneous ascending auction commonly used in spectrum auctions worked well because of its excellent price discovery, which allowed bidders to arbitrage across substitutes, and allowed the piecing together of complementary packages. In electricity markets, poor designs have arisen from attempts to oversimplify the problem: the use of a single settlement, rather than multiple settlements; the use of single-part bids, rather than three-part bids; and the use of a single price, rather than locational marginal prices.

Inadequate attention was paid to the implications of market power. Designs that magnified, rather than mitigated, the problem of market power were common: forbidding forward trades in California and creating easily manipulated capacity markets in New England. In all markets, the absence of a demand response has contributed enormously to the vulnerability of the markets to the exercise of market power. It is essential that at least some electricity consumers have an incentive to see and respond to the real-time price of electricity.

In the United States, the initial approach of the Federal Energy Regulatory Commission (FERC) to electricity restructuring was based in part on wishful thinking, as captured by the phrase: "Let a thousand flowers bloom." In the absence of a clear consensus on how markets should be restructured, FERC encouraged each market to come up with its own design. Although the design efforts were extensive, most of the designs had serious flaws.

The flawed designs should not be a surprise. Electricity market design is complex and not well understood. But perhaps the biggest impediment to good design is the fact that the designs were created by committees of stakeholders. Just as one should be hesitant to fly on an airplane designed by a committee of stakeholders, one should be hesitant to trust electricity designs that are built from consensus among interested parties. Like airplanes, electricity designs should be largely the work of experts focused solely on the objectives of the market. The compromise inherent in the design should reflect the optimum balance among competing design objectives, rather than a distributional compromise among those with conflicting interests.

Faced with more weeds than flowers, FERC has abandoned its hands-off approach in favor of a standard market design. The standard design will accomplish two things. It will eliminate as many flaws as possible, and it will encourage trade across regional markets. Perhaps the greatest motivation for restructuring is the recognition that there potentially are large gains in both efficiency and reliability from expanding the geographic scope of electricity markets. These gains are lost to the extent that incompatible market designs prevent trades across regional markets.

In this paper, I examine some principles of market design. I illustrate the principles with examples of both good and bad designs. I discuss one of the main design challenges—dealing with market power. I then examine FERC's choice of a standard market design.

2. As simple as possible, but not simpler

Good market design begins with the objective. In the case of electricity, the objective is the efficient and reliable

production of electricity to satisfy demand. This objective includes both short-run elements (least-cost daily production) and long-run elements (induce efficient investment decisions).

The second critical step in market design is to understand the preferences and constraints of market participants. This requires a knowledge of the basic economics facing suppliers and demanders. Without this knowledge, one cannot begin to assess what design elements are essential and what are unnecessary. Good knowledge of the underlying environment enables the designer to keep the design as simple as possible, while avoiding the common error of oversimplifying the design. A number of examples illustrate this principle.

The importance of price discovery

As an example of a successful design, consider the simultaneous ascending auction used to auction telecommunications spectrum licenses in the U.S. and elsewhere. In these auctions, the government is auctioning a number of related licenses to firms that hope to provide wireless communications services. The primary objective of the government is to get the licenses in the hands of those companies that value them the most. The difficulty is that some of the licenses are complements and some are substitutes, and the structure of preferences may differ across the bidders.

In the simultaneous ascending auction, all the licenses are on the block at the same time. The auction proceeds in a number of rounds with prices on each license ascending in response to bids. In each round, bidders can bid on any of the licenses. The auction ends when no bidder is willing to bid higher on any of the licenses.

The success of this design stems from its solving the most important problem bidders face. Bidders must decide what package of licenses they should acquire, given their valuations and budget constraints. The virtue of this design is its excellent price discovery. The bidders get to observe prices throughout the auction and this information gradually improves as the auction proceeds. Good price information allows the bidders to arbitrage across substitute licenses. Price information also allows the bidders to piece together complementary packages of licenses.

Most designs face some basic tradeoffs and the simultaneous ascending auction is no exception. The source of price discovery—the revelation of all bids throughout the process—can also promote undesirable collusive bidding strategies. For example, the bidding can be used as a mechanism to decide who should win what, without the determination resting on price competition. In some auctions, where collusive behavior is thought to be

an important issue it makes sense to limit this possibility by, for example, not disclosing bidder identities.

Another important tradeoff in the simultaneous ascending auction is the decision to allow bidding on individual licenses only, rather than allow the bidders to bid directly for packages of licenses. Restricting the bidding to individual licenses greatly simplifies the design, but it exposes bidders to risks if there are strong complementarities. Then a bidder may win some of what it needs, without realizing the complementarities it hoped for. If licenses are primarily substitutes or the structure of complements differs little across bidders, then allowing bidding only on individual licenses is a useful simplification. But if complementarities are large and different across bidders then such a simplification is likely to impair efficiency. Letting bidders bid directly on packages of licenses is better.

Day-ahead bids should be financially binding

Similar tradeoffs are seen in electricity market design. Here the tendency to settle on designs that are too simple has been a common mistake.

One example is the single settlement system adopted in New England and elsewhere. In a single-settlement system, participants submit bids one day ahead. The bids are used to establish the supply schedule for the next day that satisfies demand at least cost subject to reliability constraints. Suppliers are paid the energy clearing price in each hour that they produce. However, the clearing price is calculated based on the bids and resources that are available in real time. Participants can change their bids and availability until the last moment without penalty. This creates an incentive for gaming. For example, a supplier may decide close to real time to withhold or raise its bid for a generator that is scheduled day-ahead. The supplier may do this to increase the clearing price that it receives. As we move closer to real-time, the system becomes less responsive as options vanish. The supply curve becomes steeper. Hence, the vulnerability to gaming near real time is great, especially if a lot of volume is riding on the real time prices.

One solution to this problem is forbidding changes. This may be effective for bid changes, but outages are often necessary and it is difficult for the regulator to distinguish between legitimate outages and those intended to raise the price. Also, generators often have good reason to change bids in response to export opportunities, revised fuel prices, or other changes.

A better solution is to make the day-ahead schedule financially binding. This is called a multiple settlement system, since there are at least two sets of prices and quantities: those in the day-ahead schedule and those at

real time. Having the day-ahead bids financially binding does two things. First, it makes the bids credible, since successful bids involve a financial commitment. This is a general principle of market design. Bids should be financially binding. Second, a multiple settlement system mitigates incentives to manipulate the real-time price. Most of the pricing and allocation is done day-ahead. The real-time market is only to price deviations from the day-ahead plans. Those scheduled day-ahead have no incentive to manipulate the real-time price. Rather their incentive is to make adjustments to bids in response to changes in their economic situation.

The multiple settlement system illustrates a fundamental point about markets. Forward markets play an important role in both reducing risk of market participants and in reducing incentives to exercise market power. Especially in markets like electricity, where the product cannot be stored and where generators have limits on how quickly they can be started and ramped up or down, both supply and demand become more inelastic as we get closer to real time. This means that markets closer to real time will be both more volatile and more vulnerable to manipulation. Hence, both buyers and sellers have an incentive to lock in prices and quantities well in advance of real time. The market design should encourage such forward transactions. Most of the market volume should occur well in advance of real time.

A fundamental flaw in the California design was a rule barring the former utilities, which had large load obligations, from purchasing power in advance of the day-ahead market. This restriction exposed the utilities to enormous risk and made the market vulnerable to manipulation. For disaster to strike, all that was needed was a period of tight supply.

Send the right price signals

Another market simplification is having a single energy clearing price in all locations. This is the approach used in New England. Although this approach would be a useful simplification in markets without any binding transmission constraints, most if not all electricity markets often have binding transmission constraints. When transmission constraints bind, more expensive generation must be dispatched to satisfy the constraint; that is, generation that is offered at a price above the energy clearing price must be accepted. This out-of-merit generation is paid its bid, creating an extra cost that is not recovered from the energy clearing price. Instead, this transmission congestion cost is paid by all load as uplift (a constant fee per MWh across all load).

The problem with this single price approach is that it sends the wrong price signal. In the transmission-

constrained regions, the cost of generating electricity is higher than in the unconstrained regions, yet all load pays the same price for electricity. Load in the constrained regions is in essence being subsidized by the load in the unconstrained regions.

Bidding incentives also are distorted by the single-price approach. Generators in the constrained region face pay-as-bid pricing, rather than uniform pricing. Hence, their incentive is to raise their bids to the point where they are just accepted, rather than to bid something closer to their true cost. The single price approach encourages certain types of gaming to create congestion. A supplier may “bid for uplift” by creating congestion through its bids and then submitting high bids for the units that are needed to relieve the congestion.

The solution is to price congestion by using locational marginal prices. Under this approach, the price at each location reflects any congestion costs or benefits. Hence, the price in congested regions is higher than the price in unconstrained regions. Generators are paid and load is charged the appropriate locational marginal price. Of course, in an unconstrained system all the prices will be the same. But as constraints appear the prices adjust to reflect congestion costs or benefits from supply or demand at the particular location.

Let the bidders express directly the underlying economics

Another simplification is the use of single-part bids. This is done in many markets, including California and New England. Suppliers submit an increasing supply schedule for each hour. The schedule is intended to reflect the marginal cost of operating the unit at various levels. The bidders cannot express directly their true preferences, which are complicated by start-up costs and no-load costs. Consider a bidder with a generating unit that has substantial start-up costs. To express the start-up cost indirectly the generator must guess how long the unit will be asked to run. For units near the margin, an accurate guess may be impossible.

The solution is to allow multi-part bids in which generators bid not only energy but also start-up and no-load costs. The generator can express directly its true preferences. Assuming the bids are truthful, the unit commitment optimization will then determine an efficient schedule, based on better information about costs. The desirability of three-part bids hinges on the system being centrally optimized, as in the Eastern US.

Three-part bidding has an additional advantage in electricity markets. Absent market power, bidders are able to—and have the incentive to—truthfully bid their costs. In contrast, with single-part bids, generators are unable to

express their true preferences. This makes market monitoring more difficult. With single-part bids, generators can justify a high energy bid, well in excess of the unit's marginal cost, by asserting that such a bid is necessary to cover start-up and no-load costs given the expectations the bidder has about how long the unit is likely to run.

3. Avoid markets with fundamental flaws

One might think that ill-conceived markets would have little chance of surviving a lengthy design process. Surely, the most glaring flaws would be identified and eliminated before the market is implemented. Unfortunately, this is not the case. The experience to date with electricity restructuring provides numerous examples of basic market flaws not only surviving the design process, but also enduring for an extended period after the flaws are identified. Typically, some group of market participants benefits from the flaws, and if the group is large enough, it can block moves to correct the problem.

The capacity markets in New England provide a good example of markets with basic flaws. When the markets opened in May 1998 there were two capacity markets, the operable capability (OpCap) market and the installed capability (ICap) market. These markets are intended to promote reliability by assuring that there is sufficient capacity in the system to cover the peak load plus the reserve margin. ICap is a monthly market and OpCap is a daily market. Any generating unit in New England can supply ICap up to its rated capacity. If the unit is operable on a particular day, it can also supply OpCap up to its rated capacity. Load is required to purchase sufficient ICap and OpCap to cover its peak load plus a reserve margin. Generators submit bids expressing their willingness to supply both products. The clearing price for each product is found as the intersection of the aggregate supply curve with the vertical demand curve. Generators with bids below the clearing price supply the capacity product and are paid the clearing price.

At first glance, this may seem like a sensible design, but a closer look reveals that it is fundamentally flawed. One basic problem is that the products are defined in such a way that they do little to promote reliability—the objective of the markets. Reliability comes from having sufficient operable resources that are sufficiently flexible to handle contingencies as they arise. ICap and OpCap have nothing to do with the responsiveness of resources, and little to do with a resource's ability to produce energy consistently and at reasonable prices. Flexibility and the ability to produce energy are rewarded in the energy and reserve markets. As capacity becomes scarce, both energy and reserves are priced higher.

Even if it was thought desirable to further reward available capacity beyond what it can receive in the energy and reserve markets, these capacity markets are a poor method of pricing capacity. Consider the OpCap market. This market is bid day-ahead and is cleared after the fact. The units that bid below the clearing price are designated as providing OpCap and those that bid above are not. However, notice that there is no difference in the costs or risks incurred by those participants who receive payment in the market and those who do not. Winners and losers are not identified until well after the "service" has been provided by both. Every participant is providing the same service, but only those designated are paid. As a result the only rational bids in the market are a bid of zero (to insure selection in the hope there is a positive price) or a bid that is an attempt to set the clearing price. The winning bidders are receiving payment for product delivered, but the losing bidders are delivering the product as well without receiving any payment.

The true supply curve then is 0 up to the capacity available. Hence, the competitive price is either zero when there is sufficient capacity or infinity when there is a shortage. Intermediate bids and prices are the result of the exercise of market power and have nothing to do with the cost of providing capacity. It is in this sense that the prices in the OpCap market are arbitrary.

The best way to think of capacity products is as options to supply energy. The difficulty with these markets is that unlike standard options where the strike price is specified in advance, here the seller of the option is able to set its strike price ex post through its bid in the energy market. Such an option is worthless to the buyer. However, in the OpCap market the buyers (load) are required to buy a fixed quantity of this worthless product, regardless how high the price is.

Such a market is manipulated easily provided some suppliers are sufficiently large and there is not too much surplus capacity. This is indeed what happened in both the OpCap and ICap markets. Whenever the operable capacity became tight, say as the result of outages, OpCap prices would become arbitrarily high. The market was eventually eliminated, although it operated for over a year. The ICap market, which suffers from similar problems, was not eliminated until August 2000, over two years after the markets began. ICap's elimination followed a period where a few suppliers set the ICap price at levels that were arbitrarily high.

It is remarkable how long these markets persisted. The flaws were first identified in a review of the markets that was done for the system operator before the markets opened [2]. The report recommended the elimination of both markets. Not surprisingly suppliers fought for these markets as an extra revenue source and were successful for

over two years in the case of ICap. Neither market was eliminated until it become obvious that the capacity markets simply were a means of transferring an arbitrary amount of money from load to suppliers. Although ICap was eliminated as a market, ICap payments to suppliers continue, but the payment is now set administratively rather than by the suppliers.

4. Understand the politics

The persistence of bad market designs is not surprising when one examines the governance structure. In New England, market rules are ratified by a committee of interested parties. It is especially difficult to make changes that would adversely impact a large and organized group of participants, such as the suppliers. This is certainly the case with the elimination of capacity markets. In such an environment, the debate is likely to focus more on issues related to the distribution of money among stakeholders, rather than issues of efficiency. For greatest impact, designers need to propose changes that will benefit all major parties. Otherwise, the changes are likely to be vetoed.

5. Address the essential complexities

The most basic complexity in electricity markets is the real-time balance of supply and demand. This is necessary because of an inability to efficiently store electricity. The task is made difficult because of uncertainties in both supply and demand. Supply shocks can be especially large and unpredictable as a result of the failure of a large generating unit or transmission line. This fact requires that the system operator have a significant quantity of flexible resources that are quickly able to ramp up or down in response to contingencies. Properly rewarding these flexible reserve resources is a major challenge of the market design.

A second complexity is a near vertical demand curve. Today there is little demand response to price. Unresponsive demand is especially a problem because supply is highly concentrated in many markets. This means that the market is vulnerable to the exercise of market power on the supply side, and that the demand side is unable to protect itself by curtailing demand in response to high prices. Moreover, the problem of market power becomes worse as we get closer to real time. The supply curve becomes steeper as options vanish closer to real time.

Bidding nonconvexities further complicate the design problem. On the supply side, there are start-up and no-load costs. Generating units have minimum run times and the ability to ramp up and down is limited. These constraints create intertemporal dependencies. A unit may have to be

started and ramped up hours before it is needed. Industrial demand may also have intertemporal dependencies, such as a plant that requires energy over several adjacent hours to complete its production process.

A fourth source of complexity are network constraints. Typically, every action impacts everyone on the network. As a result, pricing must reflect the network constraints to induce efficient behavior.

6. Understand incentives

Almost all electricity markets are structured as uniform-price auctions. The auction determines a clearing price that balances supply and demand, and all supply offers below the clearing-price are accepted. This form of auction is vulnerable to the exercise of market power. Large suppliers have an incentive to reduce the quantity that they offer in order to get a higher price on their remaining quantity [1].

The problem of market power is exacerbated in current electricity markets by three factors. First, the response of demand to price is often negligible. The suppliers face a nearly vertical demand curve, because end consumers neither see nor pay the wholesale price. This greatly increases the profitability of price manipulation, since the higher price does not lead to a curtailment in demand. Second, the variability of supply and demand means that there inevitably are times when supply is tight, especially as we move closer to real time. Even small suppliers can influence price in these circumstances. Third, the markets sometimes suffer from excessive concentration in areas with transmission constraints. The transmission constraints may imply that demand must be met from just a few suppliers. These dominant suppliers often have both the incentive and ability to manipulate prices.

Forward contracts play an important role in mitigating the incentive to exercise market power. A supplier that has sold 100% of its generation forward has no interest in manipulating the spot price. And indeed both suppliers and demanders have an interest in forward contracts, since the contracts reduce the risks of both parties. The forward contracts are financial contracts enabling the parties to lock in acceptable prices. Importantly, the contracts do not require physical delivery. Both parties can make adjustments in response to changed circumstances. For example, a supplier may decide to satisfy its load obligation through spot purchases, whenever the spot price is below the supplier's costs. Ideally, the spot price only is used to price deviations from positions taken in forward markets. Less is riding on the spot price and so the incentive to manipulate it is greatly reduced.

The pricing rule affects behavior

It is sometimes argued that switching from uniform pricing to pay-as-bid pricing will reduce the incentive to exercise market power. See [3] for a more detailed treatment of this issue. With uniform pricing, the bidder recognizes the possibility that its bid may set the clearing price it receives for the entire quantity it wins. This gives the supplier an incentive to overstate its costs and for the overstatement to increase in the quantity bid. With pay-as-bid pricing, the supplier's bid on the last MW does not impact what the supplier is paid on the earlier MWs; hence, the incentive to overstate does not grow with quantity. However, it does not follow that the exercise of market power will be less under pay-as-bid pricing.

What we do know is that the bidding incentives depend fundamentally on the pricing rule.

With pay-as-bid pricing, the bidder's incentive is to bid as close to the clearing price as possible. Indeed, the pay-as-bid auction may be renamed "Guess the Clearing Price." The pay-as-bid auction rewards those that can best guess the clearing price. Typically, this favors larger companies that can spend more on forecasting, and are more likely to set the clearing price as a result of their size.

In sharp contrast, uniform pricing favors the smaller companies (or those with small unhedged positions going into the market). With uniform pricing, the big suppliers make room for the smaller rivals. The small suppliers are able to free-ride on the exercise of market power by the large suppliers.

Thus, the exercise of market power with pay-as-bid pricing by favoring larger bidders will tend to encourage consolidation and discourage entry; whereas the exercise of market power with uniform pricing encourages entry and reduces concentration. As a result, the market may evolve to more competitive structures under uniform pricing. This self-correcting feature of uniform pricing is especially valuable in markets like electricity that are repeated regularly.

Pay-as-bid pricing has a further disadvantage stemming from the "Guess the Clearing Price" incentives. The guesses may be wrong. In particular, a low-cost unit may guess that the clearing price will be higher than it turns out to be. As a result, more expensive units will be scheduled in its place. Hence, the dispatch is likely less efficient under pay-as-bid pricing at least in the absence of market power.

7. Demand response is essential

We argued above that demand response is important in mitigating market power. Indeed, some demand response is essential in avoiding market failure in situations of

inadequate supply. In such circumstances, if the demand curve is truly vertical it is not possible to determine a price that clears the market. A demand response solves this practical problem. But equally important demand response can greatly improve reserve markets. Demand response is a form of reserves, providing the flexibility that the ISO needs to balance supply and demand in real time. Even a little demand response will go a long way in improving the energy and reserve markets. The experience in other markets suggests that people respond to prices if they see and feel them. Real time metering is an essential prerequisite. This should happen quickly for industry and more slowly for residential customers.

One of the arguments against demand response is that consumers of electricity should not be exposed to the high volatility of the spot market. But this argument misses the important role of forward contracts. A consumer can have a forward contract that locks in a price for 100 percent of its anticipated electricity needs. The consumer is exposed to the spot market only to the extent that its needs change. The consumer, however, does have the incentive to curtail its demand when prices are high and expand its demand when prices are low. Market efficiency only requires that the consumer see and pay the spot price on the margin. By shifting demand away from high-priced hours to low-priced hours, the consumer can reduce both the electricity bill and its variance.

Arguably most of the efficiency gains from restructuring will come from the demand side. It is the demand side that has received almost no attention; whereas, even before restructuring, substantial effort went into achieving least-cost dispatch among suppliers. Unfortunately, nearly all markets have paid too little attention to the demand side, and many markets began by effectively killing retail competition. This is unfortunate, since the much needed innovation on demand management systems and contracts is likely to come from retail competition.

The common political response to restructuring was to guarantee that consumers would benefit. The restructuring legislation often would include a rate cut for consumers in the early years. For example in Massachusetts, consumers were given a 15 percent rate cut. Companies wishing to supply to retail customers had to offer consumers rates that beat this 15 percent rate cut to get them to switch from the former utilities. No companies could beat this regulated rate and make a profit, so retail competition was nonexistent.

The same thing happened in California, but with an important difference. In Massachusetts, the former utilities were protected from price risk. If the cost of supplying load was greater than the regulated rates, then the former utilities would be compensated for actual costs through a

charge that would be added to consumers' bills over a period of years. In this way retail competition was eliminated but without any price risk for the former utilities. In contrast, in California, the utilities' retail revenues were fixed at the regulated rates. The utilities did not receive any additional compensation in the event wholesale prices exceeded the regulated rates. This together with the fact that the utilities were not allowed to purchase power in forward contracts meant that the utilities were exposed to enormous price risk. When wholesale prices rose, it was simply a matter of time before the utilities ran out of money and the market collapsed.

8. FERC Standard Market Design

In testimony before the US Senate Committee on Energy and Natural Resources (September 17, 2002), Pat Wood, Chairman of FERC, conceded, "The wholesale power market today has many of the worst features of both regulated and competitive markets, and few of the benefits of either." FERC's "let a thousand flowers bloom" approach to market design had failed. In its place, FERC has proposed a Standard Market Design to be applied throughout the United States.

The move to a standard market design is compelling. Stable and consistent market rules will encourage investment and promote efficient trade across market boundaries. Most importantly, the final design will be the result of extensive study and debate, and will be informed by the past successes and failures of actual markets around the world. Ideally, the design will include all key ingredients of success and avoid the flaws that have led to market failures.

Below I summarize key elements of the design as proposed in FERC's Notice of Proposed Rule Making on Standard Market Design (July 31, 2002). I then comment on one of the features that is not adequately described.

Forward bilateral transactions will be central in the market, accounting for vast majority of the volume. These trades will be supported by (1) a forward market for congestion revenue rights to hedge transmission congestion risk, and (2) day-ahead and real-time markets for energy and reserves. The day-ahead and real-time markets are financially-binding security-constrained market-clearing auctions for energy and reserves. To manage congestion and send the right price signals, locational marginal prices are used. Units are paid the higher of market-clearing prices or bid cost recovery over the day. This is effectively the system that has been in operation in PJM and New York, and is planned for several other regions.

The markets are run by an Independent Transmission Provider (ITP), which typically will be the Regional Transmission Organization (RTO). Currently there are four RTOs covering the US and much of Canada. The ITP must be independent of any stakeholders. Only the ITP can file with FERC. Stakeholders provide input to the ITP, but cannot force decisions. This is much stronger independence than we have had in current markets.

Means for market power mitigation and monitoring are included in the rule. A bid cap of \$1000 applies to energy and reserve markets. More stringent price caps are applied in situations of market power, say where a unit must run for reliability. Automatic mitigation procedures, which impose additional limits on bids based on past behavior and the unit's characteristics, may be used as well.

Transmission planning is a central activity of the ITP. Transmission expansions are paid for by those that benefit from the transmission, rather than where the transmission is located. Access to the grid is rationed by price and on a nondiscriminatory basis. Additional access charges cannot be added when the transaction crosses market borders.

Demand response is recognized as critical to both reliability and the mitigation of market power. Demand response is put on equal footing with supply and transmission expansion plans.

Resource adequacy is a planning function

Capacity requirements, which appear as ICap markets in New England, PJM, and New York, are eliminated. Instead, the rule treats resource adequacy as largely a planning function of the ITP. Supply and demand are forecast and then the ITP assures that it has sufficient resources including demand response to meet demand and the reserve margin, which FERC sets at 12 percent. Load Serving Entities (LSEs) are assigned their share of demand and must have lined up sufficient supply and demand response to cover their peak demand plus the reserve margin. If an LSE is short, then it faces a stiff penalty at times the resources prove inadequate and/or its customers are shut off.

This approach is problematic for two reasons. First, it is unclear what it means to have sufficient resources lined up. One implementation would be similar to ICap markets and fraught with all the ICap problems. Second, shutting off the customers of an LSE that is short provides strange incentives, since doing so punishes the LSE's customers, not the LSE.

Since FERC recognizes that resource adequacy is a planning function, FERC should look for a solution that directly solves the problem at the planning level. Hence, if the ITP finds looking forward two to five years that

resources are inadequate, then the ITP should hold an auction to procure the extra resources that are needed. The auction would effectively subsidize the new additions an amount just sufficient to get the needed resources built.

With this approach resource adequacy could be more sophisticated than simply peak load plus 12 percent. The planning process could identify particular resources that are under provided in the market. For example, in New England, peaking units and demand response appear to be under provided in the market (the incremental value of these resources to the ITP is greater than the profits the resources can generate). The ITP could have auctions for each under-provided resource.

There are at least two practical difficulties with this approach. First, planning is difficult. Planned resources may or may not appear. There is likely to be disagreement about what resources are under provided. The second problem is that by only subsidizing new resources, there may be an incentive to move existing resources out of the market. For example, a peaking unit might be moved from one RTO to secure the subsidy in another, creating a greater need for peaking units in the original RTO. This problem is solved by subsidizing all peaking units, new or existing, at the same level. Although the cost of the subsidy would be greater in this case, it is still much less than it would be under ICap, which treats all MWs as equivalent. In practice, reliable, fast, and flexible resources help much more with reliability, than unreliable, slow, and inflexible resources. There is no reason to compensate all MWs in the same way, and certainly not by a pricing mechanism that is manipulated easily.

A planning approach to resource adequacy, with targeted subsidies for under-provided resources, appears far superior to an approach derived from the failed ICap markets.

9. Conclusion

Good market design begins with a thorough understanding of the market participants, their incentives, and the economic problem that the market is trying to solve. With this understanding in place, a good design follows almost from common sense. Certainly, many of the fatal flaws in actual electricity markets become obvious when the problem is analyzed in the right way. Why then do these flaws so often appear and persist?

One explanation is that common sense is scarce. The design problems actually are much trickier than meets the eye. There are many ways to look at the problem and only in hindsight does the problem become obvious. This is especially true in electricity markets, where the markets necessary are highly complex and many of the design problems involve serious challenges. The optimal design

of reserve markets is a good example. Even after years of study and experience these markets are not well understood.

A second explanation for the appearance and persistence of design flaws has to do with the design process. In the case of electricity this process has most often been design by a committee of interested parties. More often than not, design proposals were motivated by special interests. The final designs involved a bargaining compromise that tended to focus on the split of gains among special interests, rather than a design that best achieved the market's objective. Fortunately, FERC recognizes and addresses this problem head-on in the proposed Standard Market Design. The design mandates features where there is a consensus about what works well, and establishes a governance structure that will avoid the obvious conflicts of interest that can undermine efficient market operation.

There are still many challenges in electricity restructuring, but FERC appears now to be exercising the leadership that will be required to meet these challenges. The California disaster, by illustrating just how high the stakes are, gave FERC the political will to act. Consumers across the country will ultimately benefit from improved electricity markets.

10. References

- [1] Ausubel, Lawrence M. and Peter Cramton (2002), "Demand Reduction and Inefficiency in Multi-Unit Auctions," Working Paper, University of Maryland.
- [2] Cramton, Peter and Robert Wilson (1998), "A Review of ISO New England's Proposed Market Rules," White Paper, Market Design Inc., September 1998.
- [3] Kahn, Alfred E. , Peter Cramton, Robert H. Porter, and Richard D. Tabors (2001), "Pricing in the California Power Exchange Electricity Market: Should California Switch from Uniform Pricing to Pay-as-Bid Pricing?" Blue Ribbon Panel Report, California Power Exchange.