POWER

PWP-077

PJM's Capacity Market in a Price-Spike World

Steven Stoft

May, 2000

This paper is part of the working papers series of the Program on Workable Energy Regulation (POWER). POWER is a program of the University of California Energy Institute, a multicampus research unit of the University of California, located on the Berkeley campus.

University of California Energy Institute
2539 Channing Way
Berkeley, California 94720-5180
www.ucei.org

PJM's Capacity Market in a Price-Spike World

Steven Stoft [1]

May 7, 2000

Abstract

PJM's market was designed to rely on a capacity market instead of price-spikes to induce an adequate supply of generation. Were it not for PJM's links to other markets, this design could help keep prices below $150/MWh without any harm to the supply side of the market, though with some damage to incentives on the demand side.

But PJM must compete for reserves and energy with many other markets that annually produce energy-price spikes that exceed 100 times average cost. These external markets can buy PJM's energy, thereby forcing it into emergency conditions. One useful but unintended safety valve has been the internal exercise of a small amount of market power which has raised energy prices in PJM, thereby reducing exports, encouraging imports and curbing emergency conditions. The result has been price spikes inside PJM that are indistinguishable from those of a market that uses price spikes to induce generation investment.

Most problematic may be the fact that the current penalty for failing to meet the capacity requirement is too small to guarantee PJM the output of its required capacity. Increasing this penalty would make PJM more effective but could result in a greater exercise of market power in the capacity market.

This paper is both an analysis of two fundamentally different market designs for securing generation adequacy and a description of some basic features of PJM's market mechanism. The PJM market evolved out of the oldest centrally dispatched power pool in the U.S. and as a consequence much of its market design is determined more by its long and complex history than by the theoretical rationales discussed below. For this reason although the rules described below are basically accurate representations of PJM's rules, the rationales and objectives that are presented may correspond poorly if at all to historical fact. Instead these should be interpreted as views, some correct and some flawed, that are currently influencing the debate over which of these two market designs is the better design.

The No-Spike Market Model

Competition aligns market price with system marginal cost. Marginal cost for the system is the marginal cost of the most expensive plant dispatched. Therefore competitive prices never

---

[1] The author would like to thank Rob Gramlich and Joe Bowring for many essential suggestions and corrections. Any remaining errors are solely the responsibility of the author. Steven Stoft is a Senior Research Fellow at the University of California Energy Institute, and is a consultant to PJM's market monitoring unit. A more complete analysis of competition between price-spike and no-spike markets can be found in *Power Economics*, available at www.stoft.com. The views expressed in this paper are solely those of the author, and do not reflect either the views of PJM or of UCEI.

rise above the marginal cost of the most expensive peaking generator. Consequently, this generator cannot cover its fixed cost in a competitive system. Variations on this argument include the fact that operating reserve requirements mean some peakers will run only once every few years. This is taken to be additional proof that perfect competition provides no hope of providing adequate generation resources.

If this were the case, the only hope for preserving both a competitive energy market and reliability would be the invention of something like a capacity market. Call this design for low prices together with the use of a capacity market the "no-spike" market model.

For the sake of concreteness, assume that the most expensive peaker has a constant marginal cost of $130/MWh. Also assume that a new peaker would have a fixed cost that when amortized would come to $6.60/MWh. How would the no-spike model work in this world? First, energy prices would never go above $130/MWh. Second the capacity market would require all load-serving entities to purchase approximately 20% more capacity than their peak load, and this requirement would be backed by a penalty of $6.60/MWh, perhaps charged on a daily basis.

It is not surprising that the threat of a penalty equal to the cost of building a new generator can induce load to buy a new generator. It is more surprising that this can eliminate price spikes, and in fact this is possible only in a perfectly competitive or thoroughly price-capped market. We will approach the question of price spikes in stages.

First assume that the no-spike market is an isolated market with no import or export possibilities, and that it is perfectly competitive. This is sufficient to eliminate price spikes provided the market does not drive the price up by offering to pay very high prices for operating reserves on the rare occasions when they are in short supply. In an isolated system such high prices cannot call forth non-existent reserves, and they are not needed to induce investment because the capacity requirement takes care of that. PJM rules conform to this observation. PJM will not raise the market price simply because its operating reserve requirements cannot be met. When operating reserves are too low, it makes emergency purchases, and there is no price cap on these purchases. But such purchases do not set, and so cannot raise, the market price.

From April 1, 1998 until April 1, 1999, no member company of PJM had market-based rates, so none could bid above their regulated "marginal cost." Consequently no member company could push the price above about $155/MWh.[2] (In the first year, the highest physical marginal cost was about $155, in the second it was about $130.) When the PJM market did exceed this level it must have been due to the market-based bidding of companies located outside of PJM.

When reserves are very short and PJM is making emergency purchases and paying perhaps $1000/MWh or more, it is quite likely that the competitive price is much higher than PJM's market price of energy. Yet as we have just seen, it does no harm to the supply side to suppress this price spike.

With the introduction of market-based rates on April 1, 1999, it became possible for PJM members to push prices above $130. But to do so a generator must bid a little of its capacity in

[2] Of course congestion could and did push prices at specific locations above this level, but this paper will only concern itself with the non-congestion component of prices. In fact, in the first year, prices were very high only when there was no congestion and the price was uniform.

at above marginal cost. Technically this is an exercise of market power, and it is the intention of PJM that its member companies not exercise market power.

There is, of course, a very cogent demand-side reason to let the price exceed \$130/MWh. In a no-spike market with a 20% excess capacity requirement, supplying peak load for a few hours per year costs far more than \$130/MWh because it requires a peak-load generator that has carrying costs of about \$50,000/year. Peak load should see this cost. But even this reason has no force in a market where retail prices do not vary in real time. (Although an incorrectly low price causes no present inefficiency, high peak-load prices now will hasten the day when demand becomes price responsive.)

So far we have discussed the behavior of a no-spike market in isolation from all other markets. Things are very different when such a system must compete with other systems that allow and even encourage prices to rise to extreme heights. When prices are high outside of PJM there is a great temptation for PJM generators to abandon the PJM market and sell their power externally. This is exactly what happens, and the result is a shortage in PJM. But as long as PJM's load-serving entities (LSEs) have purchased the required amount of capacity, there is no serious problem. Any generator that has been sold in the capacity market and then exports its power can have its power recalled by PJM. This was done many times last summer on a number of occasions (probably about 15).

Although as long as LSEs have purchased the required capacity, PJM will continue to have access to the energy it needs. But, when external price spikes are combined with internal market power, the no-spike system begins importing the external price spikes. The details of this are discussed below.

External markets can cause a second, and potentially more severe, problem in a no-spike market. If the penalty for failing to purchase the required capacity is too low, the no-spike system can lose its capacity to competing systems just when it needs it most. Although the installed-capacity market (I-Cap market) together with PJM's price spikes have proven quite sufficient to encourage new generation, the penalty is insufficient to guarantee this capacity will be available when needed. So far PJM's LSEs have always been able to purchase the required capacity in the I-Cap market, though it is not obvious why that has been the case. This coming summer may well produce days on which PJM finds its capacity "exported" at exactly the time it is most needed. The penalty of the capacity market is designed so that it is better to build generation than to suffer the penalty every day without end. But the penalty is imposed only one day at a time which means it is only \$158/MWday. This is not a high price to pay for the opportunity to earn, say, \$2000/MWh for 16 on-peak hours.

In summary, it can be said that in isolation a non-spike market model with a correctly set capacity requirement can induce the efficient level of generation investment without the need for price spikes. But that when placed in trading contact with systems that rely on price-spikes, the non-spike market will import these price spikes through arbitrage, unless very strong measures are taken to prevent this. The surprising result is that it would operate as well or better than a true competitive market, provided two conditions are satisfied:

1. The no-spike system has no trading partners that allow price spikes.
2. Load in the no-spike system is not metered in real time.

The current PJM market very nearly satisfies condition 2 but not condition 1. When exposed to price spike markets, the non-spike market loses some efficacy on the supply side, but works better on the demand side. Before analyzing the mechanism of this failure it is worth dispelling certain myths that surround its economic competitor, the price-spike market model.

The Price-Spike Model: Regulation Through the Back Door

Competition and low elasticity imply that prices will spike when supply is short. If the supply of generation is inadequate, these spikes will be high enough to induce more investment in generation. Conversely, if there is too much installed capacity, the price spikes will be too low to induce new investment. If prices are not capped, the market will be competitive and will consequently induce an optimal level of installed capacity and consequently an optimal level of reliability.

If this argument were the whole story, then a competitive energy market and optimal reliability could be attained by closing the I-Cap market and removing any price cap. This argument is always accompanied by the assertion that the I-Cap requirement is a regulatory approach to generation adequacy while the price-spike approach is a purely market-based approach.

But something is missing: the required operating reserve margin. There can be no doubt that if the required reserve margin is 12% then the equilibrium level of installed capacity will be greater than if the required reserve margin is 2%. Since price-spike markets do have reserve requirements and since these requirements determine the level of installed capacity, it ***cannot*** be said that the optimal level of installed capacity has been determined by a purely market-driven process devoid of regulatory input. Competitive markets simply do not determine the required reserve margin.

But how does an increase in reserve margin cause an increase in the level of installed capacity? There is no direct or mechanical connection. The process does in fact work through price spikes, just as described above. All that was missing was the fact that the price spikes are largely determined by two non-market regulatory decisions: the operating reserve margin and the way it is enforced,. The process works as follows. If the reserve margin is 2% then, when available reserves are 4%, the system operator will ***not*** bid up the price of reserves so the price of energy will remain low. But, if the reserve margin is 12% and only 4% reserves are available, then the system operator will offer to pay a high price for energy, say, $1000/MWh. Consequently price spikes occur a little sooner and last a little longer in a system with a 12% reserve margin than they do in a system with a 2% reserve margin, given the same installed capacity. The result of this is that the system with the high reserve margin will induce more capacity to be built.

The regulatory setting of a reserve margin determines the demand level at which the price spike occurs. Similarly, the regulatory rule for pricing reserve purchases determines the shape of the price spike. If the rule is sharp, e.g. pay up to $9000 if reserves fall below 6% and pay nothing if they are above 6%, then the price spike will be very sharp. That is, it will be high and of short duration. If the rule is a gradual one, the price spike is lower and broader. It is neither the height nor the width that is required to induce investment in generation; it is the area, i.e. the profit, of the price spike that matters.

From this discussion, several conclusions may be drawn.

1. Whether price spikes produce optimal generation investment depends on a regulatory decision: the setting of the reserve requirement.

2. The height of the price spike depends on the rule for purchasing reserves.

3. Even with a fairly low price cap, if the reserve requirement is set high enough, the optimal amount of installed capacity will be induced.

Just as it is possible to set the wrong capacity requirement, so it is possible to set the wrong operating reserve requirement. It should also be noted that the operating reserve approach has its effect through a price-spike that is determined to a considerable extent in any given year by random variables such as the weather. This randomness cannot help but slow the process by which the market converges to its long-run equilibrium. Thus there is no automatic answer to the question of how best to induce generation adequacy. There is at this time no market solution to the reliability question because we do not yet have a market for reliability. A competitive market for energy is only guaranteed to supply energy efficiently. It has no ability to supply reliability efficiently.

But this leads to the true advantage of the price spike approach. If it is correctly implemented, it will supply energy efficiently. The benefit of the price-spike approach is that it can send the correct price signals to load. The capacity requirement approach cannot do that. While loads are currently incapable of receiving these signals, the signals are still important because they will encourage load to secure the means to respond to them. This will not only improve the efficiency of the market by improving the load shape, but more importantly it will provide the best long-run antidote to market power: demand elasticity. Lack of price-responsiveness on the part of load is the single most important cause of current market power problems.

Point 3 above has an important implication. Because of their extraordinarily low demand elasticities, current electricity markets are frequently susceptible to the (sometimes extreme) exercise of market power. This can most easily be limited with a price or bid cap, and as point 3 indicates this need not prevent the market from sending the correct investment signal to generators. Recognizing that there is no pure market approach to securing the right level of installed capacity, we proceed with a closer look at the problems of the I-Cap market.

Exporting Capacity

As noted above, an isolated no-spike market can operate optimally on the supply side. But when another market of the price-spike variety exists nearby, the no-spike market has few defenses and the price-spike market takes advantage of this vulnerability. In order to understand this interaction, it is useful at first to assume that the no-spike market continues to act as it was designed to act. That is, the internal energy price does not exceed $130/MWh.

Consider what happens when, because of either a genuine shortage of reserves or because of market power, the price in the price-spike market reaches $1000/MWh. At this price, generators in the no-spike market would be delighted to export power instead of selling their product for a mere $130/MWh. But of course there is the I-Cap market enticing them to sell internally. How strong is that enticement? At most the load-serving entities will pay up to $6.60/MWh to avoid the fine for being short of capacity. In fact, in PJM, this fine is doubled on days when the system is short of capacity. Because the fine is imposed on a daily basis it amounts to $6.60x24 or $158/MWday. Doubling this and rounding off we will say the fine is $320/MWday. This is how much a local LSE would pay for access to this generation capacity.

But how much would the generator gain by exporting? Clearly it would gain at least $870/MWh when the external price is $1000/MWh. So in half an hour it would make more than enough to cover the opportunity cost of leaving its home market. This has ***not*** happened in PJM. Generators have not withheld their capacity from the PJM market in order to export power to high-priced neighboring markets. This phenomenon is termed "exporting capacity," although the capacity is not actually sold to anyone outside of PJM except in the sense that it allows power from these generators to be sold as firm power. Such exports of power are termed "capacity-backed" and under PJM's rules they cannot be recalled even under emergency conditions. Capacity-backed exports are common, but they have not occurred to the extent that I-Cap market has gone short of capacity.

Because hot weather is highly correlated geographically, PJM experiences the most exports on the days that it most needs the energy internally. If generators do short the I-Cap market, it will almost certainly be on a day when PJM most needs the capacity. In other words it is likely that the capacity market (absent any market power) works quite well except when it is most needed.

PJM has only one defense against this reliability-threatening development. It can declare an emergency and then purchase emergency power. These purchases are not allowed to set the PJM market price so they do not do anything to stem the flow of exports. In fact if these purchases were extensive and regular they would increase the flow of exports. This is because such purchases constitute extra demand in the external market which will raise the price in that market. In fact because supply in these markets will tend to be very inelastic when it is short, these purchases could have a significant impact on the external price. Thus the more PJM imports through emergency purchases, the more attractive capacity-backed exporting becomes.

Fortunately for PJM, there is another mechanism that tends to prevent capacity-backed exports. This mechanism is not to be found in the PJM rules, and it is officially discouraged. In fact it is perhaps the main target of PJM's market monitoring unit. That mechanism is market power.

Importing Price Spikes

Perhaps PJM's major defense against emergency conditions is market power rather than the I-Cap market. When PJM runs out of supply and the energy price is $130/MWh, this presents a impressive opportunity for market power. By pricing the last MW of supply at $1000/MWh, a generator can raise the market price to that level and then receive that price on all of its power not sold in the forward market. In fact suppliers do not only bid their last MW at this price, they typically bid a sequence of small generators at increasingly high prices. But we are not concerned here with the benefit to the generator of exercising market power or its exact strategy for doing so.

The point of interest is that, by raising the market price well above the $130 allowed by PJM in the absence of market power, these generators make exporting far less attractive. In fact sometimes the PJM price is raised well above the price in the surrounding markets, so there is no incentive to export.

The complete mechanism works as follows. Say a neighboring market runs short of supply to the point where prices increase dramatically. But assume that at first PJM is only moderately short so that it does not use up all of its spare capacity. With a low price in PJM and a high price in the external market, PJM generators, even those that have been sold in PJM's I-Cap market, export to the neighbor. This uses up all of PJM's spare capacity. Now PJM's rule is that it can recall these exports because they are not capacity backed. But before doing so PJM must take every single MW of power that has been bid in day ahead. Thus if one MW was bid in at \$1000/MWh, PJM must buy that MW and thereby set the PJM market price to \$1000/MWh. It can then recall the exports (bilateral sales to the neighbor), but it must pay them the PJM market price, which is now \$1000/MWh. This is PJM's only recourse when it first runs short. Even if there are potential imports offering to sell into PJM at say \$300/MWh, these are not allowed to set the PJM market price, so PJM cannot purchase them until it accepts the internal bid at \$1000/MWh. Then it can recall exports or buy imports, but in either case it must pay \$1000/MWh.

By this mechanism, the external price is imported into the PJM market. Ironically PJM's rules, forbidding price setting by those who did not bid in the day-ahead market, can act to increase the price spike as it is imported. Of course internal market power is also to blame for the high PJM prices. But without this market power, the PJM market might have found itself in much more serious trouble. Without any internal exercise of market power, that is without bids in the day-ahead market that exceeded \$130/MWh, PJM could not have had a market price above \$130/MWh. This would have deprived it of a considerable amount of imported power. It would also have made exporting far more attractive.[3] It would have dramatically raised the opportunity cost of selling into the I-Cap market. Consequently PJM might well have experienced capacity-backed exporting last summer. These exports would have been non-recallable. Combining this with a dramatic reduction of imports would have forced PJM to rely much more heavily on out-of-market emergency purchases.

PJM needs only a very small exercise of market power. It needs this only because it has deprived itself of any other mechanism by which a shortage of supply relative to demand can raise the market price. Once PJM institutes normal market-clearing mechanisms (some of which have had to await the implementation of a binding day-ahead market on June 1, 2000), market power will no longer serve any purpose. If penalties in the capacity market are increased to the point where capacity withholding and capacity-backed power export is no longer a threat, this would also eliminate the benefit of market power. But this will not eliminate market power, and normal market-clearing mechanisms will still be desirable to help combat its exercise and reduce its impact.

How Severe Are the Problems with the Capacity Market?

In fact the capacity market has not caused any dramatic problems. It was not able to hold energy prices down to \$130/MWh or anything close to that, so load was signaled that price responsiveness could pay off. (Of course demand cannot yet respond in any case.) It provided insufficient signals for investment in generation; prices in the capacity market averaged less than half of what was needed. But the energy market more than made up the difference, and the

[3] It is not clear that PJM needed prices that were nearly as high as some of those it experienced in order to secure the necessary imports or prevent too great a level of export.

combined signal was more than adequate. It probably did not provide a sufficient penalty to prevent capacity-backed exports on crucial days, but market power pushed PJM prices so high that exporting lost much of the urgency it would otherwise have had, so the I-Cap market never ran short of supply.

Perhaps the most serious problem created by the capacity market is that it has provided yet another arena for the exercise of market power. There is no proof in this case, but three circumstances point in that direction. First, the market appears to have extremely low elasticity of demand. (The only possible exception is that Active Load Management (ALM) can be used to reduce the capacity requirement of LSEs. Although this has fluctuated by only about 300 MW out of a total required capacity demand of 60,000 MW, it may be somewhat price sensitive.) Second there were informal complaints of market power before the daily market was added. Third, it is very difficult to explain why the price in winter months has not fallen to zero and in fact was not dramatically lower than during the previous summer. During the winter, the opportunity cost of not-exporting capacity is almost nil, and there was always an excess of capacity relative to demand.

Can the I-Cap Market Guarantee that Required Capacity Will Be Available When Needed?

Although many different fixes could be proposed for PJM's I-Cap market, one in particular is natural to consider. As discussed above, the central problem of the current arrangement is that the penalty for being deficient in capacity is applied one day at time. The obvious solution to this problem is to make the requirement and the penalty backing it an annual requirement with an annual penalty. (Of course there is no theoretical support for this time period; one could as easily specify a week, a month or two years.) In other words an LSE with a requirement of 10 MW of capacity would have to purchase that capacity for the entire year and if found to be short by one MW on one day of the requirement year, it would be penalized as if it had been short by one MW for the entire year. In other words it would be penalized $58,000 instead of $158.

The effect of the annual penalty would be to keep the energy of an I-Cap generator inside PJM. There are two parts to this mechanism. First, if the energy is sold out of PJM, PJM can, in an emergency, recall it. Second, the I-Cap generator will have a harder time selling out of PJM because it cannot sell firm power. The market for non-firm power is much thinner and commands a lower price.

In order to analyze the I-Cap market, it is necessary to understand the opportunity cost associated with selling into that market. The opportunity that is relinquished by such a sale is essentially the opportunity to sell power out of PJM. The value of that lost opportunity is the profit that would be made on outside sales minus the profit that would be made on inside sales of energy.

Using a generator with a marginal cost of $30/MWh, PJM's hourly average LMPs (market prices) show that this generator, if it had run whenever the price was above this level, would have earned a short-run profit of about $77,000/MW. If that same generator had been able to avail itself of the highest neighboring on-peak day-ahead prices including the PJM price, it would have earned roughly $70,000/MW more, or $147,000.[4] Both of these numbers need to be

[4] From private data on neighboring day-ahead markets during the summer of 1999.

adjusted. It appears that a generator that had sold capacity into the I-Cap market could have made another $7000/MW by selling out of the PJM market on days when there was essentially no chance of recall. There is also the possibility of making money on a day with a recall despite having to pay liquidated damages for part of the day, but this is risky business. For the sake of argument let us adjust the inside earnings up to $90,000/MW. The external earnings need to be adjusted down because of difficulties with trading in the external market. This adjustment is very uncertain, but, again for the sake of concreteness, let us deduct $27,000/MW from this value. This gives an external profit of $120,000. The difference between external and internal is then $30,000/MW. This estimate is extremely rough not only because of the uncertainties just mentioned but because the relevant values are next summer's values, not last summer's.

What does this tell us about the price in the I-Cap market? It says the I-Cap price should not be much below $30,000. Risk aversion should push it a little below this opportunity cost but not dramatically lower. We also know that the price cannot go above $58,000, because no purchaser of I-Cap need pay more than that penalty. If the market is fully competitive, the price should be a little below $30,000, while if there is sufficient market power it will be $58,000.

Between their profits in the I-Cap market and in the energy market, generators are pretty much guaranteed to earn at least the profits available in the external price spike and perhaps more due to market power. It should not be surprising that generators are guaranteed the profits of the external price spike. This is the standard result of effective arbitrage. Of course if the transmission link between PJM and outside markets is too weak, this guarantee may fail. But this has already been at least partially taken into account in the above estimate, and indications are that this will not be too large a problem. Of more concern is the possibility that PJM generators would earn more in the PJM market than in the external market because of the extra profit from market power in the I-Cap market.

Market Power in the I-Cap Market

An example will be useful in describing the potential for market power in the PJM market. Assume there are 10 players with approximately equal supplies of generating capacity. Assume further that their total generating capacity is 62 GW and that the required capacity is only 60 GW, so that there are 2 GW of excess capacity in this market. Normally such a market would not be considered in serious danger of the exercise of market power. Its HHI is only 1000. Now let us further assume that 9 of the 10 players have no interest in raising the price because their I-Cap requirements exactly match their available capacities. In a Cournot market this would reduce the markup by a factor of ten.

This leaves us with one player, which we assume has a capacity of 6 GW, one third of which is not required by the market. Further assume that the opportunity cost with selling into the I-Cap market is $30,000/MWh for the reasons listed above. As a consequence of this, the competitive price in the I-Cap market is $30,000/MW.

How should this one player bid in the I-Cap market? It has essentially two choices. It could bid competitively and offer its capacity at a price of $30,000, its opportunity cost. In this case it should sell almost all of its capacity (97% of all capacity offered in this market will be sold). In this case it can expect to net almost 6 GW x $30,000/MW = $180,000,000 (or 97% of this to be more precise). On the other hand it could exercise market power and offer all of its capacity at

$57,000/MW. In this case it will be sure of selling two thirds of it for a profit of 4 GW x $57,0000 or $228,000,000. Clearly exercising market power is the rational alternative.

This example is useful for demonstrating more than the extreme vulnerability of this market to the exercise of market power. If the price cap (imposed indirectly through the penalty) were reduced to $44,000, the effect would be dramatic. This would not simply reduce the profits from market power to $176,000,000, but, because this figure is lower than what can be earned by bidding competitively, it would completely eliminate the exercise of market power. One danger in high prices caps (or no price caps at all) is that they actually increase the frequency with which market power is exercised.

This market behaves badly for exactly the reason that most current electricity markets have problems with market power. It has zero demand elasticity. Fortunately this condition comes to an abrupt end at a price of $58,000. But up to that point there is little to curb even a small player.

Pros and Cons of the Annual Penalty

The rationale for securing generation adequacy through a capacity market has changed at different stages in the market's development. The following seem to have been major objectives (though not necessarily those of PJM):

1. Securing the right level of investment in the generation market.
2. Paving the way for the elimination of price-spikes in the energy market.
3. Reducing the potential for the exercise of market power.
4. Securing the availability of generating capacity in real time.

The combined first and fourth objectives seems to have been the original justification for PJM's installed capacity requirement. It is now clear that an annual penalty would probably be effective at achieving this goal while the current design is not. However it also appears that PJM's policy of paying for reserves is not that different from what would be expected in a price-spike market (there seems to be no formal policy on this point). This makes perfect sense because PJM has been, like it or not, part of a price-spike market. The result is that PJM's price spikes have been very close to what would be expected from such a market, so price-spikes have replaced the I-Cap market as the major investment incentive. In any case, this does not seem to be the driving force behind the present desire to revise the I-Cap market.

The second objective, preventing price spikes in PJM, is not possible without fundamentally altering the PJM market. The first step in the alteration would be an annual penalty in the capacity market.

The third objective was perhaps based on the notion that any price above about $130 would have been a clear signal of market power, thus making detection easy. Competition with outside markets makes this test unreliable. Instead, the capacity market has simply introduced another market in which the exercise of market power appears to be problematic.

The last objective is the one driving the so-far informal search for a fix, but so far it has not gotten much play. There is a legitimate concern that PJM will run short of available generating capacity this summer so that it will be forced to shed load. Basically there are two ways to approach this problem. Buy energy or options on energy in a forward market, or wait and buy energy in real time. The use of a high I-Cap penalty is an attempt to buy call options on 60 GW of power on a year-around basis. This should provide greater security because it is possible that in real time the required energy will simply not be available. There is no particular reason that forward purchases should be any cheaper, in fact the extreme risk aversion of buyers probably means this is not the case. There is also no reason to expect the I-Cap market to perform well as a competitive market.

To summarize, the desirability of an annual penalty seems to rest on the outcome of the following tradeoff. Is the added security worth the cost imposed by the risk premium and the exercise of market power in the I-Cap market? This presumes that the annual penalty is in fact great enough to secure the I-Cap requirement. If the perceived profit available in the external market exceeds the perceived profit in the PJM market by more than $58,000, this may not be the case. But it seems likely that an annual penalty would in fact do the trick and keep the required capacity available to PJM. The main uncertainties regarding the desirability of this approach seem to be in the amount and value of the extra security that would be provided and the cost of market power exercised in the capacity market.