ATLANTIC
ECONOMICS

MEMORANDUM

DATE: February 6, 2025

TO: CAISO Price Formation Enhancements Working Group

FROM: Mike Cadwalader

RE: Comments on Fast-Start Pricing Presentations by MISO and PJM Staff and Their Independent Market Monitors

---

This memo contains my comments on some issues raised in the fast-start pricing ("FSP") presentations made by MISO,[1] PJM,[2] and their respective independent market monitors ("IMMs")[3] at the Dec. 19, 2024, and Jan. 16, 2025, meetings of the Price Formation Enhancements Working Group. While those presentations addressed a variety of issues related to FSP, I will focus on the following issues that seem to be most relevant to the working group's current FSP discussions:

- Assessing different approaches for setting LMPs.
- Relaxing integer constraints for fast-start generator ("FSG") commitment variables.
- Permitting offers from offline FSGs to set prices.

[1] MISO, Fast Start Pricing/ELMP at MISO, available at https://stakeholdercenter.caiso.com/InitiativeDocuments/Presentation-Fast%20Start-Pricing-ELMP-at-MISO-Dec-19-2024.pdf.

[2] Vijay Shah, Fast Start Pricing, available at https://stakeholdercenter.caiso.com/InitiativeDocuments/Presentation-PJM-Fast-Start-Education-Jan-16-2025.pdf.

[3] Carrie Milton, IMM Review of ELMP in MISO ("MISO IMM Presentation"), available at https://stakeholdercenter.caiso.com/InitiativeDocuments/Presentation-MISO-Independent-Market-Monitor-Review-of-ELMP-Dec-19-2024.pdf; Catherine Tyler, Fast Start Pricing in PJM ("PJM IMM Presentation"), available at https://stakeholdercenter.caiso.com/InitiativeDocuments/Presentation-PJM-Fast-Start-Pricing-Jan-16-2025.pdf.

10 Tower Office Park, Suite 408, Woburn, MA 01801
T: 781-281-0553 C: 617-875-7723 mdc@AtlanticEconomics.com

ASSESSING DIFFERENT APPROACHES FOR SETTING LMPS

Lost opportunity costs ("LOCs") arise when a generator is not dispatched to produce energy using capacity that it had offered at a price that is less than the LMP. In such cases, the generator would have increased its profits if that capacity had been dispatched to produce energy. LOCs provide an incentive for inefficient behavior because they encourage generators to disregard dispatch instructions in order to increase their profits. Alternatively, they may encourage generators to modify their offers to ensure that they are dispatched. That is also inefficient, because offers that do not reflect actual costs impair the ISO's ability to determine the most efficient dispatch to meet load.

Bid cost recovery ("BCR") payments arise when a generator is dispatched to produce energy, even though its offer to generate that amount of energy exceeds the product of the LMP and the amount of energy it was directed to produce. This can also lead to inefficiency, because the total amount paid to these generators reflects their offer to produce the amount of energy they were dispatched to produce, rather than the LMP. That gives them an incentive to increase their offers, thereby increasing the amount they are paid. But if such a generator overestimates the amount that it can offer while still being dispatched to operate, another resource that is actually more expensive, but whose bid is nevertheless lower, may be dispatched to operate in its stead, which is inefficient.

Ideally, the procedure used to calculate the LMP will seek to avoid both LOCs and BCR payments. When there are no nonconvexities—i.e., if the Pmin for every generator is zero, its start-up cost is zero, and its minimum load cost ("MLC") is less than or equal to the cost of increasing its output above Pmin—it should always be possible to find a single price that will clear the market. In other words, in those cases, it should always be possible to find an LMP that will set the amount of energy that suppliers were willing to provide at that LMP equal to the amount that buyers were willing to purchase at that LMP.[4] In that case, there will not be any generators that incur LOCs, nor will there be any generators that require BCR payments.

Of course, there are nonconvexities, and markets that CAISO and every other ISO or RTO operates must accommodate them. Even when there are nonconvexities, it is possible that a single price will clear the market. In those cases, there is an LMP at which no generator will have an LOC and at which no generator will require a BCR payment. If the implementation of FSP would cause some generators to incur LOCs, or

[4] For simplicity, I will ignore locational price differences in this discussion.

to require BCR payments, in situations where a single price can clear the market, the use of FSP would introduce inefficiency.

Example 1: A Single LMP Clears the Market

Consider the following example.

- The FSG has a Pmin of 150 MW, and a Pmax of 200 MW.
    - The FSG submits a start-up bid of $2000, which is allocated over a minimum up time ("MUT") of one hour, and an MLC of $7000/hour, so its commitment cost (i.e., its bid cost of starting and operating at Pmin for one hour) is $9000, or $60/MWh at its Pmin of 150 MW.
    - The FSG submits an offer of $80/MWh for its dispatchable segment.
- There are two other non-fast-start generators available: G1 and G2.
    - Each has a Pmin of 0 MW and a Pmax of 500 MW.
    - G1 offers all of its energy at $35/MWh, while G2 offers all of its energy at $85/MWh.
- There are 675 MW of load.

Table 1 shows the least-cost dispatch to meet load.[5] G1 is dispatched to operate at Pmax, producing 500 MWh. The FSG is dispatched above Pmin, producing 175 MW, while G2 does not operate, as it is more expensive than the FSG.

Table 1: Least-Cost Dispatch for Example 1

	Capacity (MW)	Offer ($/MWh)	Schedule (MW)	Bid Cost ($)
G1	500	$ 35.00	500	$ 17,500
G2	500	$ 85.00	-	$ -
FSG:				
Pmin	150		150	$ 9,000
Inc	50	$ 80.00	25	$ 2,000
Total			**675**	**$ 28,500**

[5] The examples in this memo will determine physical dispatch instructions and LMPs for the hour as a whole. Other complicating factors, such as operating reserves or FRP, will be disregarded (with the exception of Example 3 in a later section, which will consider operating reserves).

In previous meetings, we discussed three different ways of producing a modified offer curve for the FSG, which would be used in the pricing dispatch in lieu of the actual offer that was submitted for that FSG. One of those methods was the constant adder approach. In this example, under that approach, the ISO would calculate an adder for the FSG of \$45/MWh,[6] which it would add to the bid that was actually submitted for the dispatchable segment of the FSG, yielding a modified offer for the FSG's full capacity of \$80/MWh + \$45/MWh = \$125/MWh. As Table 2 shows, using that modified offer in the pricing dispatch would produce an LMP of \$85/MWh, since G2 would be on the margin in the pricing dispatch, even though it was not actually dispatched to operate.

Table 2: Pricing Dispatch for Example 1, Using the Constant Adder Approach to Determine the Modified Offer for the FSG

	Capacity (MW)	Actual Offer (\$/MWh)	Modified Offer (\$/MWh)	Pricing Schedule (MW)
G1	500	\$ 35.00	\$ 35.00	500
G2	500	\$ 85.00	\$ 85.00	175
FSG:				
Pmin	150		\$ 125.00	-
Inc	50	\$ 80.00	\$ 125.00	-
Total				**675**

That LMP would not clear the market, because the FSG has an LOC. The FSG was actually dispatched to operate at 175 MW, so it will be paid \$85/MWh × 175 MWh = \$14,875 for that energy. Since its as-offered cost of producing that amount of energy is \$9000 + \$80/MWh × 25 MWh = \$11,000, it will realize \$3875 in profit. But it would prefer to have been dispatched to operate at Pmax, as it would then realize an additional (\$85/MWh – \$80/MWh) × 25 MWh = \$125 in profit on the 25 MW that it offered at a price that is less than the LMP.

In contrast, if either the adjusted constant adder approach or the minimum average cost approach is used, the modified offer curve would be \$60/MWh for the FSG's Pmin block and \$80/MWh for the incremental 50 MW. The LMP would be \$80/MWh, as Table 3 shows, since the FSG's dispatchable segment would be on the margin.

[6] The adder is equal to the commitment cost divided by Pmax, or (\$2000 + \$7000) / 200 MWh = \$45/MWh.

Table 3: Pricing Dispatch for Example 1, Using the Minimum Average Cost Approach to Determine the Modified Offer for the FSG

	Capacity (MW)	Actual Offer ($/MWh)	Modified Offer ($/MWh)	Pricing Schedule (MW)
G1	500	$ 35.00	$ 35.00	500
G2	500	$ 85.00	$ 85.00	-
FSG:				
Pmin	150		$ 60.00	150
Inc	50	$ 80.00	$ 80.00	25
Total				**675**

Thus, in this example, there is a price, $80/MWh, that clears the market and produces efficient incentives. If the modified offer curve used in FSP is constructed using the adjusted constant adder approach or the minimum average cost approach, the LMP will be that price. Similarly, if FSP is not used in this example, the LMP would be the market-clearing price of $80/MWh. In contrast, if the modified offer curve used in FSP is constructed using the constant adder approach, the LMP that results would not produce efficient incentives, as it would result in LOCs.[7]

Example 2: A Single LMP Does Not Clear the Market

However, nonconvexities will often lead to situations in which no single LMP clears the market. Consequently, there is no LMP that will eliminate both LOCs and BCR payments. Higher LMPs will lead to LOCs, while lower LMPs will lead to BCR payments, but one or the other (or both) must occur. In such cases, one should consider the

[7] The modified offer curves that are constructed using the adjusted constant adder approach and the minimum average cost approach are identical in this example, but in other examples, in which the offer for the FSG to produce energy above Pmin is less than the average cost of producing energy at Pmin, the adjusted constant adder approach will produce a different modified offer curve than the minimum average cost approach produces. In such cases, the adjusted constant adder may result in LMPs that do not clear the market.

My memo commenting on CAISO's analysis of fast-start pricing that was presented at the working group's April 8, 2024 meeting contains an example that illustrates this point. In that memo's Example 2, the adjusted constant adder approach resulted in an LMP of $95/MWh. In that example, just as in this one, an LMP of $80/MWh cleared the market. An LMP of $95/MWh would have caused the FSG to incur an LOC because it would have increased its profits if it had been dispatched to operate at its Pmax of 200 MW. See Mike Cadwalader, Comments on Analysis of Fast-Start Pricing, available at https://stakeholdercenter.caiso.com/InitiativeDocuments/Mike%20Cadwalader%20Comments%20Price%20Formation%20Enhancements%20Working%20Group%20Session%2016%20-%20April%208,%202024.pdf.

impact on both LOCs and BCR payments when assessing different approaches for calcualting LMPs.

With this in mind, consider the following example. It is almost identical to Example 1, except that:

- G2 offers all of its energy at $70/MWh, instead of $85/MWh.
- There are 625 MW of load, instead of 675 MW.

Table 4 shows the least-cost dispatch to meet load.

Table 4: Least-Cost Dispatch for Example 2

	Capacity (MW)	Offer ($/MWh)	Schedule (MW)	Bid Cost ($)
G1	500	$ 35.00	475	$ 16,625
G2	500	$ 70.00	-	$ -
FSG:				
Pmin	150		150	$ 9,000
Inc	50	$ 80.00	-	$ -
Total			**625**	**$ 25,625**

The FSG is dispatched to operate at its Pmin of 150 MW. As a result, it is necessary to reduce G1's output to 475 MW, below its Pmax of 500 MW.[8]

There is no LMP that clears the market in this example. On one hand, the FSG's average cost to operate at 150 MW is $60/MWh, so the LMP must be set at $60/MWh or higher to avoid the need to make a BCR payment to the FSG. But an LMP above $35/MWh will cause G1 to incur an LOC, as in that case, it would be better off if it could generate 500 MWh rather than the 425 MWh it has been dispatched to produce. Since the LMP cannot be both less than or equal to $35/MWh and greater than or equal to $60/MWh, any approach to setting the price will result in either LOCs or BCR payments.

In this example, adopting FSP, and using the minimum average cost approach to produce the modified offer for the FSG, would result in an LMP of $60/MWh. As Table 5 shows, that is because the modified offer for the FSG's Pmin block is $60/MWh, and that block is on the margin in the pricing dispatch.

[8] An alternative is not to commit the FSG and to dispatch G1 to operate at 500 MW while G2 operates at 125 MW, but the bid cost of that dispatch is 500 MWh × $35/MWh + 125 MWh × $70/MWh = $26,250, so it is more costly than the dispatch shown in Table 4.

Table 5: Pricing Dispatch for Example 2, Using the Minimum Average Cost Approach to Determine the Modified Offer for the FSG

	Capacity (MW)	Actual Offer ($/MWh)	Modified Offer ($/MWh)	Pricing Schedule (MW)
G1	500	$ 35.00	$ 35.00	500
G2	500	$ 70.00	$ 70.00	-
FSG:				
Pmin	150		$ 60.00	125
Inc	50	$ 80.00	$ 80.00	-
Total				**625**

While it would not be necessary to make any BCR payments if the LMP is $60/MWh, there would be an LOC for G1. The LMP exceeds its $35/MWh offer by $25/MWh, so it would make an additional of 25 MWh × $25/MWh = $625 in profit if it had been dispatched to operate at its Pmax of 500 MW. If the ISO does not make payments to G1 to compensate it for this lost opportunity, G1 would have an incentive to engage in behavior that would increase its profits by ensuring that it operates at Pmax. For example, it could increase its Pmin to 500 MW. Alternatively, the ISO could make LOC payments to G1 to eliminate this incentive, but this would produce an incentive for G1 to reduce its offers below its incremental cost of operation in order to increase its LOC payments.

However, there are also adverse consequences associated with not adopting FSP. In this example, if FSP is not implemented, the LMP would be set at G1's offer, which is $35/MWh, because G1 would be dispatched up to meet an increment of load. At this LMP, there would not be any LOCs. But there would be a $3750 BCR payment to the FSG, as it only receives $35/MWh × 150 MWh = $5250 in energy payments, based on the $35/MWh LMP, while its offer to produce 150 MWh was $9000. Effectively, the generators are paid two different prices: G1 receives the LMP of $35/MWh, while the FSG is paid $9000 / 150 MWh = $60/MWh for each MWh that it generates. The large gap between these two prices provides a strong incentive for G1 to modify its offers so that it will be paid something closer to the price that the FSG receives.

Thus, if G1 was to increase its offer from $35/MWh to $60/MWh, it would continue to be dispatched to produce 475 MWh, but if FSP is not used, the LMP would increase from $35/MWh to $60/MWh. In response to a question that I posed, the PJM IMM suggested that market monitoring could apply mitigation to ensure that G1 does not increase its offer in this manner, but that presumes that this increase in the offer is an attempt to exercise market power, and it is far from clear that it is. G1 is not withholding any energy; it continues to be dispatched at 475 MW even with the increased offer.

Instead, G1 is simply modifying its bidding strategy in an attempt to ensure that it is paid the same price that the ISO is paying to the FSG.

More generally, the first defense against market power should not be market monitoring. Instead, the first defense should be a market that is designed that, as often as possible, gives generators an incentive to submit offers that reflect their actual cost structure. If generators that are not attempting to withhold nevertheless have an incentive to submit offers that do not reflect their actual costs, it becomes much more difficult for the ISO to distinguish between offers that actually intend to drive up prices by withholding capacity and offers that are not submitted with that intent. Consquently, this may make it easier to disguise attempts to exercise market power, and it may also increase the likelihood that mitigation is applied when it is not appropriate. But when there is a gap between the prices that are paid to some generators and the prices that are paid to other generators at the same location and time, there is a incentive for inframarginal generators like G1 to increase their offers even when they do not intend to withhold energy. And when there is a large gap between those prices—as in this example—that incentive is stronger.

In addition to making it more difficult to identify attempts to withhold capacity, a market design that gives generators an incentive to bid in this manner can also introduce inefficency, because inframarginal generators may incorrectly assess how high they can bid. In this example, if G1 was to increase its offer above $70/MWh in the mistaken belief that it would continue to be dispatched, G2 would be directed to operate instead of G1. That would be inefficient, since G2 is actually less expensive than G1, but that is a consequence of providing incentives for generators to submit offers that do not reflect their cost.

There is no perfect solution in cases where the presence of nonconvexities means that no single price will clear the market. Instead, the decision as to whether to implement FSP or not, and which version of FSP to implement if FSP is implemented, should be based on an attempt to minimize the incentives for inefficiency that will result, with the understanding that they cannot be avoided completely.

Relaxing Integer Constraints for FSG Commitment Variables

Instead of constructing a modified offer curve for use in the pricing dispatch to determine the LMP, the approach used by both MISO and PJM relaxes integer constraints that reflect whether an FSG is committed or not. In the physical dispatch, the commitment variable for an FSG must be equal to either 0 (indicating that the unit is not committed) or 1 (indicating that the unit is committed). If the unit is committed, then the full start-up cost associated with that unit is incurred; also, the unit must incur its MLC while operating at no less than its Pmin for its MUT. But if this constraint is relaxed for an FSG in the pricing dispatch, then the FSG can be fractionally committed. Thus, if the commitment variable for an FSG is equal to 0.5, only half of the start-up cost would be committed, the unit would be required to incur half of its MLC, and it could be dispatched to operate at no less than half of its actual Pmin, but no more than half of its actual Pmax.

Examples 1 and 2

In Examples 1 and 2 above, the LMPs that result from relaxing the integer constraint for the commitment variable for FSGs are $80/MWh in Example 1 and $60/MWh in Example 2.

- In Example 1, if the constraint requiring the commitment variable to be an integer is relaxed, the FSG would still be fully committed, and the value of the commitment variable would be 1. Therefore, an increment of load would be met by increasing output by the FSG's dispatchable segment, and the LMP is the $80/MWh offer for that segment.
- In Example 2, as Table 6 shows, if the constraint requiring the commitment variable to be an integer is relaxed, the FSG would be five-sixths committed. The FSG's modified Pmin is 0.8333 × 150 MW = 125 MW, and its modified Pmax is 0.8333 × 200 MW = 166.67 MW. An additional MW of load would be met by increasing the value of the FSG's commitment variable to 0.84. That would increase the FSG's modified Pmin to 0.84 × 150 MW = 126 MW, thereby permitting the FSG to produce 126 MW, rather than 125 MW, using its Pmin block. Since the commitment cost is $9000, the cost of increasing the value of the commitment varaible from 0.8333 to 0.84 is (0.84 – 0.8333) × $9000 = $60, so the LMP is $60/MWh.

Table 6: Pricing Dispatch for Example 2, Relaxing the Integer Constraint for the FSG Commitment Variable

	Capacity (MW)	Offer ($/MWh)	Mod. Capacity (MW)	Pricing Schedule (MW)
G1	500	$ 35.00	500.00	500.00
G2	500	$ 70.00	500.00	-
FSG:				
Pmin	150		125.00	125.00
Inc	50	$ 80.00	41.67	-
Total				**625.00**

These are the same LMPs that are produced in these examples by using the modified offer curve (if the minimum average cost approach is used to construct the offer curve). In general, in energy-only examples, I think that relaxing the integer constraint for FSG commitment variables will yield the same LMPs as using a modified offer curve that is constructed using the minimum average cost approach. However, the two methods may produce different results in examples that include an operating reserve ("OR") requirement, as Example 3 will show.

Example 3: Energy and Operating Reserves

Example 3 is similar to Examples 1 and 2. The differences are:

- G1's output is offered at $42/MWh, while G2's output is offered at $80/MWh.
- There are 575 MW of load.
- There is a 30 MW OR requirement, which can be met by any generator, using any capacity that is not scheduled to generate energy.

Table 7 shows one of the least-cost schedules to meet load and the OR requirement. The FSG is dispatched to operate at its Pmin of 150 MW. As a result, it is necessary to reduce G1's output to 425 MW, below its Pmax of 500 MW.[9] G1 is scheduled to provide 30 MW of OR. However, either G2 or the FSG could have been scheduled to provide some or all of that OR, instead of G1, without increasing the bid cost of the schedule.

Table 7: Least-Cost Dispatch for Example 3

	Capacity (MW)	Offer ($/MWh)	Energy Schedule (MW)	OR Schedule (MW)	Bid Cost ($)
G1	500	$ 42.00	425	30	$ 17,850
G2	500	$ 80.00	-	-	$ -
FSG:					
Pmin	150		150	-	$ 9,000
Inc	50	$ 80.00	-	-	$ -
Total			**575**	**30**	**$ 26,850**

If the minimum average cost approach is used to construct the modified offer used in the pricing dispatch, we obtain the dispatch shown in Table 8.

[9] An alternative is not to commit the FSG and to dispatch G1 to operate at 500 MW while G2 operates at 75 MW and provides 30 MW of OR, but the bid cost of that dispatch is 500 MWh × $42/MWh + 75 MWh × $80/MWh = $27,000, so it is more costly than the dispatch shown in Table 7.

Table 8: Pricing Dispatch for Example 3, Using the Minimum Average Cost Approach to Determine the Modified Offer for the FSG

	Capacity (MW)	Actual Offer ($/MWh)	Modified Offer ($/MWh)	Pricing Pass Energy Sch. (MW)	Pricing Pass OR Sch. (MW)
G1	500	$ 42.00	$ 42.00	500	-
G2	500	$ 80.00	$ 80.00	-	-
FSG:					
Pmin	150		$ 60.00	75	30
Inc	50	$ 80.00	$ 80.00	-	-
Total				**575**	**30**

The LMP is $60/MWh because the FSG's Pmin block is on the margin in the pricing dispatch, and the modified offer for that block is $60/MWh. The price of OR is $0/MWh because an increase in the OR requirement could be met by scheduling any of the generators to provide more OR, without affecting the bid cost of meeting load and the OR requirement.[10]

In contrast, if the integer constraint for the FSG commitment variable is relaxed, we obtain the dispatch shown in Table 9. The FSG is 52.5% committed, so its modified Pmin is 0.525 × 150 MW = 78.75 MW and its modified Pmax is 0.525 × 200 MW = 105 MW. Thus, it can be scheduled to produce 78.75 MW of energy using its Pmin block and 26.25 MW of OR using its dispatchable segment.

[10] The PJM IMM indicated concern that the market rules in PJM may inflate OR prices because generators are not permitted to provide OR using their Pmin blocks. (PJM IMM Presentation at 40.) This concern appears to be valid to me. In this example, suppose that G2 is not available. Then, since the capacities of the remaining units—G1 and the FSG—sum to 700 MW, and load is 575 MW, 125 MW is available to provide OR in the physical dispatch, as Table 7 shows. As Table 8 shows, if the modified offer is used, G1 would be dispatched up to its Pmax of 500 MW, so all 125 MW of the capacity that is available to provide OR in the pricing dispatch is provided by the FSG.

However, the FSG cannot provide more than 50 MW of OR in the physical dispatch, since it cannot operate below its Pmin of 150 MW. Thus, this approach lets the FSG provide OR in the pricing dispatch using 75 MW of capacity that cannot actually provide OR in the physical dispatch. This 75 MW on the FSG offsets 75 MW of capacity on G1 that actually can provide OR in the physical dispatch, but cannot provide it in the pricing dispatch because G1 has been dispatched up to its Pmax. Therefore, implementing a rule that would limit the FSG to providing no more than 50 MW of OR in the pricing dispatch could lead to increased OR prices, potentially reaching shortage pricing levels, even though there is actually more than enough capacity available to meet OR requirements, since such a rule would cause the total amount of capacity that is able to provide OR in the pricing dispatch to be less than the 125 MW that is actually able to provide OR in the physical dispatch.

Table 9: Pricing Dispatch for Example 3, Relaxing the Integer Constraint for the FSG Commitment Variable

	Capacity (MW)	Offer ($/MWh)	Mod. Capacity (MW)	Pricing Pass Energy Sch. (MW)	Pricing Pass OR Sch. (MW)
G1	500	$ 42.00	500.0	496.25	3.75
G2	500	$ 80.00	500.0	-	-
FSG:					
Pmin	150		78.75	78.75	-
Inc	50	$ 80.00	26.25	-	26.25
Total				**575.00**	**30.00**

The LMP that results is $55.50/MWh, while the price of OR is $13.50/MWh. These prices are determined as follows:

- An additional 1 MW of load would be met by increasing the value of the commitment variable from 0.525 to 0.53, which would increase the FSG's modified Pmin to 0.53 × 150 MW = 79.5 MW and its modified Pmax to 0.53 × 200 MW = 106 MW. This permits the FSG to increase the output of the Pmin block by 0.75 MW, from 78.75 MW to 79.5 MW; the remaining 0.25 MW is produced by increasing G1's output from 496.25 MW to 496.5 MW. The cost of meeting this additional 1 MWh of load is (0.53 – 0.525) × $9000 + 0.25 × $42/MWh = $45/MWh + $10.50/MWh = $55.50/MWh, so the LMP is $55.50/MWh.

- An increase in 1 MW in the OR requirement would also be met by increasing the value of the commitment variable from 0.525 to 0.53. That would permit the Pmin block of the FSG to generate an additional 0.75 MW and to provide an additional 0.25 MW of OR. G1 would reduce its output by 0.75 MW to offset the increase in the FSG's output, thereby permitting it to provide another 0.75 MW of OR. Thus, the cost of a 1 MW increase in the OR requirement is (0.53 – 0.525) × $9000 – 0.75 × $42/MWh = $45/MWh – $31.50/MWh = $13.50/MWh, so the OR price is $13.50/MWh.

Neither of these combinations of energy and OR prices will clear the market.

- The combination of a $60/MWh LMP and a $0/MWh OR price will not result in BCR payments, but it will cause G1 to incur LOCs of ($60/MWh – $42/MWh) × (500 MW – 425 MW) = $1350, as G1 would prefer to be dispatched to produce 500 MW of energy rather than 425 MW, since the LMP exceeds its energy offer.

- The combination of a $55.50/MWh LMP and a $13.50/MWh OR price would necessitate BCR payments of $9000 – 150 MWh × $55.50/MWh = $675 for the FSG. It would also lead to LOCs of $13.50/MWh × (500 MW – 455 MW) = $607.50 for G1, since G1 would prefer to be scheduled to produce either energy or OR (it does not care which, since it makes a profit of $13.50/MWh on either) with the 45 MW of its capacity that is not scheduled in the DAM. Finally, it would lead to LOCs of 500 MW × $13.50/MWh = $6750 for G2, since G2 would like for all 500 MW of its capacity to be scheduled to provide OR.

In this example, there is much more capacity available in the physical dispatch to meet operating reserve requirements than is needed to meet those requirements. But, if prices are calculated by relaxing the integer constraint for the FSG commitment variable, we obtain positive OR prices because the pricing dispatch only commits the fraction of the FSG that is needed to meet load and OR requirements. An increase in the OR requirement leads to the need to increase the fraction of the FSG that is committed. The resutling increase in cost sets the OR price.

In the physical dispatch schedule that was shown in Table 7, G1 was scheduled to provide 30 MW of OR, but that was an arbitrary choice. G2, or the FSG, or any combination of the generators could have been scheduled to provide OR. If the price of OR is positive, as in this example if the integer constraint on the FSG commitment variable is relaxed, and there are no costs associated with providing OR, each generator will want to be scheduled to provide it. That will give each generator an incentive to try to manipulate the dispatch to increase the likelihood that it is the lucky generator that is scheduled to provide OR. That may reduce the efficiency of the dispatch. The alternative, which would eliminate this incentive, is to make LOC payments, but those could be quite large, as this example shows.

Permitting Offers from Offline FSGs to Set Prices

One design question that has arisen in past working group meetings is whether offers submitted by FSGs that are offline should be considered when setting prices under FSP. The MISO IMM expressed its view that this is efficient only when the FSG can be started quickly to address a shortage, and it would be economic to start the FSG to address the shortage.[11] To assess this claim, consider Example 4.

Example 4

In this example, assume there are three generators.

- G1 has a Pmin of 0 MW, a Pmax of 500 MW, and offers energy at $35/MWh. It is not an FSG.

[11] MISO IMM Presentation at 6.

- G2 has a Pmin of 0 MW, a Pmax of 100 MW, and offers energy at $500/MWh. It is not an FSG.
- The FSG does not have a dispatchable segment. Its Pmin and Pmax are both 100 MW. Its commitment cost is $6000 for a one-hour MUT, so the cost of its energy is $60/MWh.
- There are 505 MW of load.

The least-cost dispatch is shown in Table 10. G1 operates at its Pmax of 500 MW, and G2 produces the remaining 5 MW that are needed to meet load. The FSG is offline.

Table 10: Least-Cost Dispatch for Example 4, When Load is 505 MW

	Pmax (MW)	Pmin (MW)	Offer ($/MWh)	Schedule (MW)	Bid Cost ($)
G1	500	-	$ 35.00	500	$ 17,500
G2	100	-	$ 500.00	5	$ 2,500
FSG (offline)	100	100	$ 60.00	-	$ -
Total				**505**	**$ 20,000**

If load is expected to remain 505 MW, it would be inefficient to start the FSG. If the FSG runs, it must operate at 100 MW, which means that G1 would be dispatched down to 405 MW. The bid cost of the resulting dispatch is 405 MW × $35/MWh + $6000 = $20,175, which exceeds the cost of the dispatch in Table 10. Therefore, the FSG should not start.

If offline FSGs are not permitted to participate in price-setting, the LMP in this case would be set by G2's offer of $500/MWh. That price does, in fact, clear the market, given the generators that are online. If the LMP is $500/MWh, neither G1 nor G2 would require a BCR payment or incur any LOCs. But the FSG would be very happy to operate at an LMP of $500/MWh. In contrast, if offers from offline FSGs are considered when setting the price, then the LMP would be the FSG's offer of $60/MWh, since it would be dispatched to produce 5 MW in the pricing dispatch. As a result, G2 would require a BCR payment, since it is actually dispatched to produce energy, and its offer is well above $60/MWh.

If load is expected to increase to 510 MW at some point in the future, then it will be efficient to start the FSG at that time. As Table 11 shows, even though it will be necessary to back the FSG down by 90 MW to accommodate the FSG's Pmin, the resulting bid cost of the dispatch is $20,350. That is less than the cost that results if G2 is dispatched to operate at 10 MW, which would be 500 MW × $35/MWh + 10 MW × $500/MWh = $22,500. Accordingly, the LMP would be set by the $60/MWh offer for the FSG.

Table 11: Least-Cost Dispatch for Example 4, When Load is 510 MW

	Pmax (MW)	Pmin (MW)	Offer ($/MWh)	Schedule (MW)	Bid Cost ($)
G1	500	-	$ 35.00	410	$ 14,350
G2	500	-	$ 500.00	-	$ -
FSG	100	100	$ 60.00	100	$ 6,000
Total				**510**	**$ 20,350**

Therefore, if offline FSGs cannot set the LMP, the LMP will be $500/MWh when the load is 505 MW, and will then drop to $60/MWh when load increases slightly and it becomes economic to start the FSG. There is an argument that this produce a useful price signal, because it is difficult to anticipate circumstances such as this, in which LMPs are very high because it is inefficient to start the FSG because load is just slightly above G1's Pmax. Consequently, the LMP may be likely to drop by the time that offline resources can respond to it. Making offline FSGs eligible to set the LMP would eliminate these transient price spikes.

On the other hand, if offers by offline FSGs are not considered in price-setting, the resulting potential for transient price spikes will encourage resources to be available to mitigate those spikes. In this example, suppose that a 5 MW battery was always available to inject energy at $100/MWh. Then it would not be efficient for G2 to run when load is 505 MW. The battery would generate 5 MW instead, and the battery's $100/MWh offer, rather than G2's $500/MWh offer, would set the LMP. In contrast, if offers from offline FSGs are permitted to set the LMP, then the LMP would be $60/MWh even when load is 505 MW. That could reduce the incentive for the battery to be available.