ATLANTIC
ECONOMICS

MEMORANDUM

DATE: January 10, 2024

TO: CAISO Price Formation Enhancements Working Group

FROM: Mike Cadwalader

RE: Comments on Presentation at Working Group Session #9

This memo contains my comments on the analysis contained in the CAISO's presentation at Price Formation Enhancements Working Group Session #9, held on December 11, and my suggestions regarding potential modifications to that analysis. In short, my recommendations and observations are:

- The two approaches that the CAISO used to construct the adjusted offer curves that are used in the pricing pass differ in two respects, so CAISO should consider adding a third approach to its analysis. This third approach is a constant adder constructed using a different method. This would facilitate the identification of the root cause of differences resulting from different approaches to constructing adjusted offer curves.

- CAISO should consider evaluating the impact that each fast-start pricing approach has on (1) the frequency with which fast-start units would recover their as-offered costs without the need for bid cost recovery payments, and (2) the lost opportunity costs incurred by resources with dispatchable capacity that is not economic to operate, but which might nevertheless be profitable to operate. These will assist in assesssing whether fast-start pricing would accomplish the objective of ensuring that LMPs reflect the commitment costs of fast-start units when those units are needed to meet load, while also assessing the impact of each approach on incentives for other resources to operate when, and only when, it is economically efficient for them to do so.

- Since fast-start units will continue to incur minimum load costs even if they continue to operate after completing their minimum up time, the CAISO should

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

consider modifying its analysis to include adjusted offers for fast-start units have completed their minimum up times, to ensure that those costs are reflected in the adjusted offers. Adjusted offers for these intervals would not include start-up costs.

- It is not surprising that most of the impact of fast-start pricing is limited to just a few hours of the day, as the units whose operation would be most likely to cause differences between LMPs calculated using fast-start pricing and LMPs calculated without fast-start pricing are units that operarte for short periods.

- Given the likeiihood that offers and bids in the day-ahead market will change to ensure that changes in day-ahead LMPs mirror the changes in real-time LMPs that the implementation of fast-start pricing would be expected to yield, analysis of the impact that fast-start pricing would have on payments by load that is based on fast-start pricing's impact on real-time LMPs is likely to be more accurate than analysis that is based on the impact that fast-start pricing would have on day-ahead LMPs, if the latter analysis does not consider the impact that implementing fast-start pricing would have on day-ahead offers and bids.

PROCEDURE USED TO DETERMINE THE CONSTANT ADDER

Under fast-start pricing (FSP), fast-start units would be modeled in the pricing pass as though they were fully dispatchable at any point between an output of zero MW and their maximum output level (Pmax). As a result, it is necessary to define an adjusted offer curve for such units, since the offers that are actually submitted for such units will not indicate the incremental cost of producing each additional MWh for output between zero MW and the minimum operating level (Pmin) for these units. Additionally, the adjusted offer curve must be non-decreasing, so that the marginal offer cannot decrease as output increases.

In order for the adjusted offer curve to incorporate all of the costs of committing and operating a fast-start unit, commitment cost—including both the start-up cost and the minimum load cost (MLC), which reflects the cost of operating at Pmin—must be reflected in the adjusted offer curve in some way. The December 11 presentation analyzed the consequences of two different methods for doing so, which I will illustrate using a simple example.

Assumptions for Illustrative Example

In the example, I assume the following:

- The fast-start unit has a Pmin of 100 MW, and offers two incremental 50 MW energy blocks, so it can be dispatched at any output level between 100 MW and 200 MW.

- The fast start unit offers a start-up cost of $2000, an MLC of $5000/hour, and submits incremental output offers of $40/MWh for the first incremental block and $80/MWh for the second incremental block.

Additionally, for the purposes of this illustration, I assume that the unit has a minimum up time (MUT) of 60 minutes, and that the commitment cost is allocated equally over the intervals within that 60-minute MUT.[1]

Given those assumptions, Table 1 illustrates how the two approaches used in the December 11 presentation—the constant adder approach and the minimum average cost approach—would be applied to determine adjusted offer curves in this example.

Table 1: Methods of Determining Adjusted Offer Curves

		Actual Offer			Constant Adder			Minimum Average Cost		
	Capacity (MW)	Offer ($/MWh)	Cost of Block	Cumul. Cost	Adj. Offer ($/MWh)	Cost of Block	Cumul. Cost	Adj. Offer ($/MWh)	Cost of Block	Cumul. Cost
Pmin	100		$ 7,000	$ 7,000	$ 75.00	$ 7,500	$ 7,500	$ 60.00	$ 6,000	$ 6,000
Inc 1	50	$ 40.00	$ 2,000	$ 9,000	$ 75.00	$ 3,750	$11,250	$ 60.00	$ 3,000	$ 9,000
Inc 2	50	$ 80.00	$ 4,000	$13,000	$ 115.00	$ 5,750	$17,000	$ 80.00	$ 4,000	$13,000

Calculation of Adjusted Offer Curves

Under the first approach analyzed in the December 11 presentation (the constant adder approach), the $7000 in commitment cost (calculated by summing the $2000 start-up cost and the MLC of $5000/hour × 1 hour = $5000) is divided by the unit's output when it operates at Pmax (which is 200 MW × 1 hour = 200 MWh) to arrive at an adder of $35/MWh. This is then added to the actual offer submitted for each output block, under the assumption that output by the Pmin block is offered at the same price as output by the first incremental block. Thus, the adjusted offers are $40/MWh + $35/MWh = $75/MWh for output anywhere between 0 MW and 150 MW, encompassing the Pmin block and the first incremental block, and $80/MWh + $35/MWh = $115/MWh for output between 150 MW and 200 MW, encompassing the second incremental block.

Under the minimum average cost approach (the second approach analyzed in the December 11 presentation), the average cost of operating at each output level is

[1] The December 11 presentation also included the results of analysis that was conducted under the assumption that commitment cost (for fast-start units that can start within 30 minutes) would be allocated evenly over a 30-minute MUT. The CAISO's analysis must make some assumption regarding how this cost would be allocated; these assumptions are reasonable. But there is no reason why commitment cost must be allocated over the MUT, and even if it is allocated over the MUT, there is no reason why it must be allocated equally to each interval within the MUT. Consequently, if the CAISO proceeds with FSP, it might elect to use an approach that makes different assumptions with respect to how commitment cost is allocated over time than the assumption that it made for the purpose of this analysis.

determined. In this example, the average cost of operating at the Pmin of 100 MW for the MUT of one hour is simply the commitment cost divided by the output at Pmin, or \$7000 / (100 MW × 1 hour) = \$70/MWh. Since 50 MW of output using the first incremental block has been offered at a price that is less than \$70/MWh, increasing output to 150 MW will reduce the average cost, to (\$7000 + \$40/MWh × 50 MW × 1 hour) / (150 MW × 1 hour) = \$9000 / 150 MWh = \$60/MWh. The second block of incremental output has been offered at \$80/MWh; since that is greater than \$60/MWh, increaing output above 150 MW will increase average cost. Thus, the minimum average cost is \$60/MWh, which is achieved at an output level of 150 MW.

Given that result, the adjusted offer curve for all output at or below 150 MW—in other words, for the Pmin and the first incremental block—is set at this minimum average cost level, \$60/MWh. The adjusted offer curve for output between 150 MW and 200 MW, encompassing the second incremental block, is equal to the \$80/MWh offer that was actually submitted, so the adder that is applied to that block is zero. (More generally, under this approach, the adjusted offer curve is identical to the offer curve that was actually submitted for any blocks exceeding the output level at which average cost is minimized.)

Comparison of Adjusted Offer Curves

There are two diffferences between these two approaches for calculating adjusted offer curves. First, as the name suggests, the constant adder applies the same adder to each block, while the minimum average cost approach does not. Additionally, the as-offered cost of operation is higher under the constant adder approach than under the minimum average cost approach. The cost of operating this fast-start unit at Pmax for one hour, using adjusted offer curves calculated using the constant adder, is \$17,000. The cost of operating this unit at Pmax for one hour using adjusted offer curves calculated using the minimum average cost approach is just \$13,000, which is identical to the \$13,000 cost of operating this unit at Pmax for one hour that is calculated using the offers that were actually submitted.

I am concerned that the presence of these two differences between these approaches will make it considerably more difficult for the CAISO to derive conclusions from its analysis comparing the impact of applying the constant adder to the impact of applying the minimum average cost approach. If LMPs are higher using the constant adder than using the minimum average cost approach, that could result from the fact that the cost of operating the fast-start unit calculated using the constant adder approach is higher than the cost of operating that unit calculated using the minimum average cost approach, but it could also be attributable to differences between how the two methods allocate commitment cost to different output blocks.

An Alternative Version of the Constant Adder

To determine how much of the impact is attributable to each of these two differences. the CAISO should consider extending its analysis to incorporate a third approach to calculating adjusted offer curves. This approach is another variant of the constant adder, but under this approach, the adder is calculated by subtracting the product of the output when operating at Pmin (100 MW × 1 hour = 100 MWh) and the offer price for the first incremental energy block ($40/MWh) from commitment cost.[2] Thus, instead of an adder of $7000 / 200 MWh = $35/MWh, the constant adder would be ($7000 – $4000) / 200 MWh = $15/MWh. Under this approach, the adjusted offers would be $40/MWh + $15/MWh = $55/MWh for the Pmin block and the first incremental block, and $80/MWh + $15/MWh = $95/MWh for the second incremental block.

Table 2: Costs Calculated Using Actual Offers, the Alternative Version of the Constant Adder, and the Minimum Average Cost Approach

		Actual Offer			Alt. Constant Adder			Minimum Average Cost		
	Capacity (MW)	Offer ($/MWh)	Cost of Block	Cumul. Cost	Adj. Offer ($/MWh)	Cost of Block	Cumul. Cost	Adj. Offer ($/MWh)	Cost of Block	Cumul. Cost
Pmin	100		$ 7,000	$ 7,000	$ 55.00	$ 5,500	$ 5,500	$ 60.00	$ 6,000	$ 6,000
Inc 1	50	$ 40.00	$ 2,000	$ 9,000	$ 55.00	$ 2,750	$ 8,250	$ 60.00	$ 3,000	$ 9,000
Inc 2	50	$ 80.00	$ 4,000	$13,000	$ 95.00	$ 4,750	$13,000	$ 80.00	$ 4,000	$13,000

As Table 2 shows, the cost of operating this fast-start unit at Pmax for one hour using this alternative version of the constant adder would be the same as the cost that would be calculated using either the offer curves that were actually submitted or adjusted offer curves that are determined using the minimum average cost approach. The CAISO could then compare results using this approach to the results that use the minimum average cost approach to determine the impact of allocating the same amount of cost in different ways.

THE IMPACT OF FSP ON THE NEED FOR BID COST RECOVERY PAYMENTS

The fundamental concept underlying FSP is that when a fast-start unit is needed to meet load at a given point in time, the LMP for that point in time ought to reflect the cost of committing and operating that generator. Consider a variant of the example given above, in which the two incremental blocks are removed, so that the fast-start unit can operate at only one output level, 100 MW, and incurs a commitment cost of $7000 if the start-up cost is amortized over the one-hour MUT. If that unit is marginal in

[2] This alternative version of the constant adder approach was described in a presentation at Working Group Session #6. Price Formation Enhancements (Oct. 12, 2023) at 26, available at: https://www.caiso.com/InitiativeDocuments/Presentation-Price-Formation-Enhancements-Oct12-2023.pdf.

the pricing pass, then it will set the LMP at its location at $7000 / (100 MW × 1 hour) = $70/MWh, and that price, if maintained for the hour, would permit recovery of its commitment cost. And if that unit is inframarginal in the pricing pass, it might recover more than its commitment cost. In either case, it would not require any bid cost recovery (BCR) payments to make up any difference between the market revenue it receives and its offers to provide the energy that it provided. For that reason, in order to assess whether the proposed approaches to implementing FSP are accomplishing the objective, it will be important to evaluate the impact that implementing FSP has on the number of fast-start units that require BCR payments.

Now, there are several reasons why implementing FSP is not likely to eliminate the need for BCR payments completely. If a fast-start unit starts, but load turns out to be lower than was forecasted or the output from solar or wind resources is higher than was forecasted when the unit was directed to start, the unit may not actually be needed to meet load, even though it was forecasted to be needed. Thus, it might not be dispatched in the pricing pass, or it might be dispatched in the pricing pass, but only for intervals covering a portion of its MUT. As a result, the LMPs it receives might not be sufficient for it to recover its cost of operation. Additionally, fast-start units with dispatchable capacity might not operate at high enough output levels to permit recovery of all of the costs included in their offers. Nevertheless, this analysis should provide some insight into whether FSP is accomplishing this objective, and extensions of this analysis are likely to be useful in the future, if the working group addresses topics such as whether commitment cost should be allocated equally to each interval within the MUT.

Additionally, as was suggested during the Market Surveillance Committee's December 18 meeting, FSP can also lead to cases when resources that would be uneconomic to operate nevertheless incur lost opportunity costs, because the LMP exceeds the marginal value of generation at a given location. In the example above, Table 1 showed that under the constant adder approach, the adjusted offer for the second incremental block of output is $115/MWh, even though the offer that was actually submitted for that block is $80/MWh. This might lead to circumstances when that unit was not directed to produce energy using that block, because the unit is dispatched based on the $115/MWh adjusted offer. But its actual marginal cost of producing additional energy using that block is just $80/MWh. Therefore, if the LMP is higher than $80/MWh, it would incur a lost opportunity cost, which would give that generator an incentive to produce energy using that block, notwithstanding the fact that it was not dispatched to operate. This can undermine the CAISO's ability to operate the system as efficiently as possible. Thus, an evaluation of the impact that each proposed approach for implementing FSP would have on lost opportunity costs would be valuable.

Adjusted Offer Costs After the MUT Has Been Completed

One of the assumptions in the analysis was that commitment cost would only be amortized through the end of the MUT. If possible, I suggest reconsidering this assumption. Certainly, if start-up cost is amortized over the MUT, it does not make sense to continue to include start-up cost when determining the adjusted offer curve for intervals that follow the conclusion of the MUT.[3] After all, start-up cost is only incurred once per start. However, fast-start units will continue to incur the MLC for each interval in whch they continue to operate following the conclusion of the MUT. If FSP is not applied during those intevals, and the MLC (stated in dollars per MWh of energy produced at Pmin) exceeds the cost of producing energy using the first incremental block, then some of the costs associated with operating the fast-start unit will be disregarded. As a result, the LMP may not fully reflect the cost of operating those units if they are needed to meet load in intervals following the conclusion of their MUT. Therefore, I recommend modifying the analysis so that it would continue to use adjusted offer curves following the conclusion of each fast-start generator's MUT. However, while the adjusted offer curves that apply during the MUT include start-up cost in the commitment cost for a fast-start unit, the adjusted offer curves that would apply to the intervals after a given fast-start unit has completed its MUT should include just the MLC in commitment cost.

Continuing with the example above: Suppose that this fast-start unit operates at 150 MW following the conclusion of its one-hour MUT. The full cost of continuing to operate this unit is its MLC plus its cost of produing energy using its first incremental block, or ($5000 + $40/MWh × 50 MW × 1 hour) / (150 MW × 1 hour) = $7000 / 150 MWh = $46.67/MWh. This exceeds the $40/MWh cost of producing energy using the first incremental block. If the fast-start unit is needed to meet load in the interval, applying FSP would permit the full cost of operating the unit to meet load in that interval to be reflected in the LMP for that interval, instead of just the $40/MWh cost of incremental output using the first incremental block of energy.

Applying FSP in this case would also provide better incentives that could facilitate more efficient operation of the system. For example, suppose that the CAISO has two alternatives available to it to meet load: it can schedule an import of 150 MW offered at $45/MWh, or it can continue to operate this fast-start unit at 150 MW after the conclusion of its MUT, at a cost of $46.67/MWh. The import is the more efficient option. But if FSP is not applied, the LMP will not be high enough to support scheduling the import. The fast-start unit would continue to operate, and the LMP would be set at

[3] Again, this assumes for the purpose of this analysis that start-up cost will be allocated over the MUT, while recognizing that whether start-up cost should be allocated over some other time period is a topic to be addressed in the future.

$40/MWh, which does not fully reflect the cost of continuing to operate this unit after it has completed its MUT.

TIME PROFILE OF IMPACT ON LMPS

The analysis contained in the report indicates that the impact of FSP is largely limited to periods between 5 a.m. and 8 a.m., and between 4 p.m. and 8 p.m.[4] While I cannot vouch for the magnitude of the effect on LMP, I am not surprised that the impact is largely limited to these periods of time. FSP will only affect LMPs when fast-start units with positive commitment costs are dispatched to operate, since those are the units whose adjusted offers will differ from the offers they actually submitted, so that using their adjusted offers in lieu of their actual offers may affect prices. Units that meet the start-up time and MUT requirements to qualify as fast-start units while still having significant start-up costs and MLCs are generally units that will operate for short periods of time, whether to meet peak demand, to respond to contingencies, or to increase or reduce output during high ramp-up or ramp-down periods when ramping constraints affect the dispatch. Therefore, most of the impact of FSP should occur during these relatively short time periods.

THE IMPACT OF FSP ON ACTIVITY IN THE DAY-AHEAD MARKET

The analysis in the December 11 presentation of the impact of FSP on payments by load assumed that all load would pay the real-time price, so it reflected just the impact of FSP on real-time prices. My understanding is that the CAISO plans to extend the analysis to reflect the impact that FSP would have on LMPs in the day-ahead market, and to use that to assess the impact that FSP would have on the amount paid by load for energy purchased in the day-ahead market.

For a given set of bids, offers, and system conditions, I would expect the impact of FSP on day-ahead LMPs to be considerably lower than its impact on real-time LMPs. First, in the day-ahead market, there are often many other options available, such as imports or energy from generators that can be committed in the day-ahead market, but which cannot start quickly enough to meet needs that are identified with less lead time. Second, and perhaps more important, is the difference in granularity. The day-ahead market determines schedules and prices on an hourly basis, which means that a full hour's ramping capabilty is available to meet changes in forecasted conditions from one hour to the next. Consequently, many of the short-term events that may lead to the need to start fast-start units in the real-time market may simply not appear in the day-ahead market. For both of these reasons, it is less likely that fast-start units with

[4] This refers to the calculations that permit units with start-up times of 60 minutes or less to qualify as fast-start units. If fast-start units are limited to units with 30-minute start-up times, FSP does not have a significant impact on LMPs in any of the hours reported.

positive commitment costs will be scheduled to operate in the day-ahead market than in the real-time market.

However, the caveat above is critical, as I would expect the implementation of FSP in the real-time market to affect bids and offers submitted to the day-ahead market. In general, the primary factor that drives prices in forward markets is expectations of spot prices, because forward positions can always be settled at spot prices. Therefore, if the application of FSP was expected to increase the average real-time LMP at a given location over a given hour by, say, \$3/MWh, I would expect that the application of FSP would increase the day-ahead LMP at that location in that hour by an amount close to \$3/MWh. Offers and bids by sellers and buyers, including virtual bidders, would change as necessary to produce that result. Therefore, I think that analysis that is based on the impact of FSP on real-time LMPs would be more likely to produce an accurate indication of how the implementation of FSP would affect the amount paid by load than analysis of the impact of FSP on day-ahead prices that does not account for this effect on offers and bids submitted to the day-ahead market.