

Opinion on Pay-for-Performance Regulation

by

James Bushnell, Member
Scott M. Harvey, Member
Benjamin F. Hobbs, Chair

Members of the Market Surveillance Committee of the California ISO

Final of March 9, 2012

1. Introduction

The Market Surveillance Committee (MSC) of the California Independent System Operator has been asked to provide an opinion on the ISO's proposal on Pay-for-Performance Regulation. This proposal has been made to respond to FERC Order 755 (October 20, 2011). In that order, the Commission required independent system operators to develop a mechanism to pay frequency regulation resources based on the actual services provided. These payments are to consist of two parts:

- a payment for capacity reserved for regulation services and
- a payment for performance based upon the amount of frequency regulation provided by resources when accurately following the automatic generator control (AGC) dispatch signals provided by the ISO.

In response to the order, the ISO has developed a proposal for regulation payment mechanism, whose most recent version was released on February 22, 2012.¹ The proposal has undergone significant evolution in response to stakeholder and other comments, and, in our view, has been highly responsive to the concerns that have been raised. In this opinion, we will comment on some issues that are inherently difficult because of the nature of frequency regulation and the specific requirements of the FERC order. We believe that the ISO's present proposal represents a largely effective response to these issues. Nonetheless, there are fundamental uncertainties about the effect of its implementation. These uncertainties mean that careful tuning of the parameters of the mechanism will be required as well as close monitoring to ensure that the desired incentives are put in place without providing opportunities for gaming and unnecessary inflation of costs to consumers.

The uncertainty arises from the fundamental difficulty that actual performance of a regulation resource (measured by 'mileage', the sum of absolute values of the movements of the resource in response to instructions) that is to receive payment is likely to differ substantially from the resource mileage that would be calculated in the ISO market optimization software when

¹ CAISO, Pay for Performance Regulation, Draft Final Proposal Addendum, February 22, 2012, www.caiso.com/Documents/Addendum-DraftFinalProposal-Pay_PerformanceRegulation.pdf

determining the bid-based price for mileage. Actual performance is governed by the control rules embodied in the AGC system, and responds to unexpected very short-term deviations in supply-demand balance. Because performance is providing service for unexpected events, the exact usage is by definition impossible to precisely anticipate in advance. This divergence between actual and assumed mileage potentially creates incentives for non-cost based bidding behavior and gaming in an attempt to inflate payments, as we will explain. The ISO's proposal attempts to minimize those incentives and their possible impacts, but we believe that the risk remains for significant distortions. Consequently, the ISO must be prepared to adjust the parameters of the regulation payment mechanism quickly if problems arise.

In the rest of this opinion, we first summarize some determinants of the value and cost of regulation services, and the need for a market to reflect these (Section 2). We then consider some market design choices and how they affect potential gaming behavior (Sections 3 and 4). In Section 5, we discuss certain parameters governing the requirements for mileage in the market software as well as the capability of resources to provide that mileage, and the need for consistency in their definition. Section 6 briefly addresses three other issues in regulation markets, including payments for imbalance energy, bid caps, and cost allocation, while Section 7 summarizes our conclusions.

2. The Value and Cost of Regulation Capacity and Mileage: General Considerations

The value of frequency regulation to the ISO system arises from the ability of regulation resources to respond quickly to changes in system supply and demand conditions to maintain system frequency as well as targeted exchanges with neighboring systems. This response consists of moving a resource (adding or subtracting supply or adjusting load) within a 10 minute interval in response to fluctuations in the supply-demand balance. Such movements may or may not restore a resource to its scheduled operating point by the end of the interval. A regulation resource can be viewed as providing several services: net imbalance energy over the interval (which may be negative); net movement in one direction over the interval; and absolute amounts of movement within the interval. The need for the first two services arises from load and supply forecast error, while the need for the third comes from the inevitable within-interval fluctuations of the supply-demand balance. Net energy is compensated for by payment at the real-time energy price. Capacity to provide net movement will be compensated for by the capacity portion of the regulation payment. Finally, absolute amounts of movement ('mileage') will be compensated for by the performance, or mileage, payment.

These services are very important to the system, but the precise amounts that are needed of each, and the extent (if any) to which more of one can substitute for less of another are uncertain. This was clear in the stakeholder and ISO discussions of the Regulation Energy Management proposal, in which no conclusion was reached concerning even the general magnitude of any value of the quicker response time of REM resources.² Under the present market design, this

² Regulation Energy Management Draft Final Proposal, California ISO, Jan. 13, 2011, http://www.caiso.com/Documents/RevisedDraftFinalProposal-RegulationEnergyManagement-Jan13_2011.pdf

greater responsiveness is not rewarded. In our opinion on the REM proposal, we stated that this value could be significant, but needed to be traded off against other resource attributes.³ In general, requirements for regulation resources are informed by operator experience and modeling analyses, such as the ISO's 20% renewables study. However, we are unaware of studies that have examined the specific amounts needed of the individual services, especially mileage; such studies are needed.

Meanwhile, the cost of supplying these services varies from resource to resource. Energy costs for traditional regulation sources, of course, depend on fuel costs and efficiency. Also, for storage units, they depend on the round-trip efficiency of the resource in storing and generating energy and the cost of charging energy. Demand-side sources incur costs in form of foregone benefits of consumption if loads are reduced to provide up-regulation. Energy bids are not considered when selecting regulation resources, and because the AGC system does not operate resources based on their cost, it is quite possible that the real-time price does not cover the as-bid energy cost of a resource.

The cost of reserving generation capacity so that it can provide a net movement over the interval is largely the opportunity cost of not being able to use that resource to provide energy or other ancillary services in that interval. The co-optimization carried out in the market software automatically calculates prices that reflect and fully cover these opportunity costs when they occur, based on the prices of the other commodities and the resource's as-bid costs of providing those commodities. Because the ISO's market software will not capture the cost of foregoing opportunities to sell energy in intervals beyond the time horizon of the software (especially for energy-limited units), some opportunity costs may need to be reflected in the resource's bids to provide capacity for regulation. FERC's order explicitly allowed such opportunity costs to be included in bids. In the case of demand-side resources, there may be costs associated with operating in a mode in which they can reduce their power use in response to AGC instructions.

Finally, costs for moving a resource up and down within an interval can arise from increased maintenance expenses, deterioration in average heat rates for traditional power plants, and the expense of charging energy for storage sources. If demand-side resources provide these services, there can also be management costs and foregone consumption benefits.

Thus, regulation provides three distinct types of services to the ISO's markets. The amounts of each that the ISO requires (especially of mileage) and acceptable tradeoffs among them are uncertain at this time, so the regulation market design will need to be flexible so that it can be adapted to changing conditions and improved operator understanding about what is needed. Furthermore, the costs of providing those services will vary considerably among resources. Since those costs can be significant, this implies that market mechanisms are desirable in order to enable resources to reveal their costs of providing these services so that they can efficiently provided. The FERC order and the ISO's proposal attempt to respond to that need. The foregoing implies that an important criterion for evaluating any proposal is the extent to which it would incent cost-based bidding.

³ Market Surveillance Committee of the California ISO, Opinion on Regulation Energy Management, Jan. 21, 2012, www.caiso.com/Documents/FinalMSCOpiniononRegulationEnergyManagement.pdf

A fundamental challenge in designing markets that would allow for cost discovery and efficient acquisition of regulation resources is the gap between how a market mechanism would choose resources to provide mileage, and how the AGC system would actually utilize those resources. Ideally, the costs, as estimated by the market software, would be a close approximation of the actual costs incurred in actual operations, with the forward market using the same objective function as in actual operations. In theory, imbalances between the schedule and actual operation should average out around zero over time so that there is no predictable difference, and they should be settled at real-time prices.

However, these conditions will not be the case with mileage provided by regulation under the ISO's proposal. There are two ways in which the scheduling and pricing of regulation mileage differs from how most commodities in the ISO's markets are scheduled and settled.

- First, the market software will schedule regulation resources and estimate their mileage contributions based on minimizing as-bid cost, subject to the bounds (4) and (5) of the Proposal; in contrast, AGC moves regulation resources to maintain area control error (ACE) within the Balancing Area ACE Limit and to preserve as much rampability as possible. As a result, the amount and distribution of AGC-determined mileage and its distribution among regulation resources is likely to deviate systematically from what the market software 'schedules.'⁴
- Second, no settlement will be made on the basis of the mileage that is implicitly 'scheduled' by the market software; the market clearing price will be calculated by the market software based on its calculation of expected mileage, but resources will be paid for their AGC-determined mileage. Thus, if the mileage selected by the market software (which determines the price) consistently differs from the AGC mileage (which determines the quantity), price may be much lower than the as-bid cost of actually dispatched mileage for some resources, while potentially being much higher than the as-bid cost of mileage on other resources that are not actually used by AGC.

These two considerations mean that there is a significant danger of strong incentives for some regulation suppliers to bid in a way that systematically deviates from costs, as we explain below.

3. One versus Two Commodity Market Clearing Constraints and Prices

The ISO's proposal involves defining two constraints for regulation in the market software: a requirement for capacity (constraint (1) on page 9 of the Feb. 22 proposal, hereafter called "the Proposal") and a requirement for mileage (constraint (3) in the Proposal). Each can be viewed as a constraint that the 'supply' of the commodity (capacity or mileage provided by regulation resources) equals or exceeds the 'demand' (the operator-imposed requirement). The objective function (preceding constraint (1) in the Proposal) includes terms representing the as-bid costs of

⁴ Thus, in this opinion we will put quote marks around 'scheduled' when referring to the amount of mileage that is calculated by the market software. This is not a binding schedule in any physical or financial sense; it is just a device for calculating a market-clearing price.

regulation capacity and mileage for the resources. The shadow prices of the constraints are the respective market clearing prices for capacity and mileage. Those shadow prices may equal the bid of the last accepted resource, or (in the case of capacity) an opportunity cost if that capacity could otherwise have generated energy or provided another type of ancillary service, or finally, a scarcity value if too few resources are available to meet the constraint.⁵

An alternative formulation that was the basis of an earlier draft proposal by the ISO would instead have established just one constraint for regulation, which would have been defined as a weighted sum of mileage and capacity. The objective function coefficient for that composite regulation product would have been based on a weighted sum of the capacity and mileage bids. Then only a single composite price would have resulted which would have been decomposed after the fact into separate capacity and mileage prices. We believe that the use of a single constraint formulation would cause several problems.

- Using a single standard mileage rate for all resources would greatly underestimate the mileage faster resources would actually provide under AGC, resulting in incentives for them to understate their mileage bid and overstate their capacity bid in order to maximize the calculated mileage price. On the other hand, slower resources would have a significant risk of under recovery of costs because of the resulting depressed capacity prices and the fact that AGC would acquire much less mileage from them than the standard mileage rate.
- If instead resource-specific mileage rates were to be used to construct bids, this could penalize faster resources by increasing their apparent cost, unless a more complex formulation were used in the objective function.
- During times of high prices for energy or non-regulation ancillary services, mileage prices might incorrectly reflect a high opportunity cost, when in reality only regulating capacity can incur such a cost.

Therefore, we find the general two constraint approach of the ISO's proposal to be preferable to a regulation market based upon a single composite constraint. However, this is not to say that a two constraint approach will necessarily be free of potential unintended consequences, as we discuss next. The particular values that are chosen for the parameters in the market optimization (page 9 in the Proposal) could significantly affect the opportunities for and consequences of gaming and the likelihood of other unintended consequences.

⁵ Although the requirement for mileage in constraint (3) cannot be more than the amount of regulation capacity that has been offered into the market weighted by their unit-specific mileage rates m_i , it is still possible for the price of mileage to be based on the scarcity price of \$55/MWh. This can occur if some of the offered regulation capacity is instead scheduled by the software to provide another commodity (energy or other ancillary services), so that the potential 'supply' of mileage is reduced.

4. Equality versus inequality formulations of the mileage supply constraints and possible gaming

In this section we discuss a key market design choice that has a profound effect on the mileage price and could lead to a gaming opportunity if the market clearing price of regulation capacity is not close to zero. The market design choice refers to the definition of a particular constraint in the market optimization. We discuss two alternatives; each strikes a different balance between the need for a positive and transparent mileage price per the FERC order and the need to avoid gaming opportunities. The approach embodied in the ISO's proposal meets the mileage pricing need, and we support it. However, this approach could result in a significant gaming opportunity if the market clearing price of regulation capacity is not close to zero. It is quite possible that with the implementation of a mileage payment, most of the costs that were previously reflected in the regulation capacity payment will be included in the mileage payment so that the price of regulation capacity will typically be close to zero, except perhaps when there are significant opportunity costs to providing regulation, but this is hard to know until we observe the new design in operation.

4.1 Alternative Constraint Formulations

A key feature of the ISO's proposal is that the supply of mileage from a resource provided in constraint (3) can be any value between:

- the MW of regulation capacity that the resource supplies to constraint (1) (defined as the capacity that can be supplied in 10 minutes) and
- a mileage multiplier m_i times that capacity, expressing the amount of mileage that the resource could provide over an hour, which would range up to 6 for a normal resource (see constraints (4) and (5) in the Proposal).

The parameter m_i could take even higher values for unconventional fast resources. For instance, a fast resource that is 20 MW in size whose ramp rate might permit, say, 10 MW of mileage per MW of regulation capacity per 10 minute period, would then be allowed to supply between $1 \cdot 20$ and $6 \cdot 10 \cdot 20$ MW of mileage to the mileage constraint (3) if all 20 MW of its capacity is selected to meet the regulation capacity constraint (1). Of course, the AGC system's use of that resource will be unaffected by the amount of mileage "scheduled", be it 20 MW or 1200 MW, as only the regulation capacity made available and its ramp rate is considered by the AGC system, not the 'scheduled' mileage. If that fast resource submitted a high offer price for its mileage, while other regulation resources scheduled for capacity have more than enough cheap mileage to meet the mileage requirement (3), then the amount of mileage that clears from the fast resource will likely be 20 MW. This may be a great understatement of the mileage the AGC system would actually use on that fast resource.

This same potential discrepancy will exist for conventional resources, although the difference would be less extreme. Based on the proposed formulation of constraint (5), even a conventional resource could be sent AGC instructions to move up to six times as much as the mileage assumed in evaluating the economics of scheduling regulation. This discrepancy potentially provides an opportunity for gaming that we explain in Section 4.2.

An alternative formulation would make constraint (4) an equality. This would mean that the mileage ‘scheduled’ in the market software would precisely equal the amount of regulation capacity schedule times the resource’s specific mileage rate m_i . As a result, high mileage resources would be ‘scheduled’ to supply more mileage than low mileage resources, which would provide a more realistic estimate of the relative amounts of mileage that AGC would utilize from the various resources.

It is our understanding that the reason that the ISO formulates constraint (4) as an inequality, thus allowing resources to be ‘scheduled’ in the economic evaluation for much less mileage than they would actually be instructed to provide, is the desirability of a readily interpretable and positive price for mileage. In particular, the inequality formulation would likely result in a shadow price for the mileage requirement (3) that reflects the price offer of the most expensive ‘scheduled’ mileage.⁶ This easily interpreted price is a seemingly straightforward implementation of the FERC order’s requirement for the mileage price to be based on mileage bids. Also, the price will be positive, unless the mileage requirement (right side of constraint (3)) is less than the capacity requirement (right side of constraint (1)), but this is very unlikely given that constraint (1) is expressed in regulating capacity per 10 minutes and the mileage requirement is expressed in mileage per hour. In that unlikely case, the minimum amount of mileage that can be provided based on constraint (5) would exceed the mileage requirement, which would cause (3) to be slack with a zero price.⁷

In contrast, using an equality formulation for (4) will, in general, result in a more complex interplay between capacity bids, mileage bids, and the prices for those two services. The mileage price could be below or above the highest accepted mileage offer price; similarly, the capacity price could also be below or above the highest accepted capacity offer price. A particular issue is that it would be more likely that the mileage constraint (3) would be slack--resulting in a zero

⁶ Under the following three distinct conditions, this would not be true.

1. A mathematical condition called degeneracy may occur (where the total requirement is exactly met by the scheduled resources, and every resource providing mileage is either at its lower (5) or upper (4) bound).
2. If constraint (5) (the lower bound) is binding for the source of the most expensive mileage supplied, then that resource will not set the price, and the mileage price will generally be lower (except possibly if degeneracy occurs). The extreme case of this occurs if the mileage requirement (right hand side of (3)) is less than the capacity requirement (right side of (1)). This could occur for the mileage requirement is based on average mileage in similar previous periods, as proposed by the ISO. Then the mileage provided each resources selected for capacity will be at its lower bound (5) and the total mileage constraint will be slack (mileage provided will be strictly greater than the mileage requirement). The result will be a zero price for mileage.
3. Third, if all resources providing mileage are at their upper bound (4), the mileage constraint may be forcing more regulation capacity to be acquired than is required by the regulation capacity constraint itself (1); then the mileage price will also reflect the cost of capacity. This can happen if the mileage required in (3) is relatively high compared to the capacity required in (1).

⁷ This assumes that $m_i \geq 1$ for all resources. If this is not the case, then constraints (4) and (5) would be inconsistent, and there would be no feasible solution; one solution would be to drop constraint (5) for resources with $m_i < 1$, and make constraint (4) an equality.

mileage price--if the regulation capacity acquired could have more than enough mileage associated with it to meet that constraint.

Under the equality formulation, in general, both the capacity (1) and mileage (3) requirements can be binding and have positive prices, although it is possible for only one or the other to be binding in any particular situation. An important problem arises if only the capacity constraint is binding, which results in a situation that we believe is likely to be contrary to FERC's intent. If the mileage requirement (right side of (3)) is relatively low, then it is quite likely that it will be below the amount of mileage that the selected capacity resources could potentially supply, and only the capacity constraint will have a positive price. In this situation, the shadow price for the capacity constraint is likely to include the marginal cost of mileage. This occurs because increasing the accepted capacity from a marginal resource will mean that more mileage will also necessarily be acquired from that unit, if (4) is an equality constraint. The mileage price itself will be zero, which is contrary to the apparent intent of the FERC order. An additional perhaps surprising result of this situation is that fast resources will be penalized in this situation, as all else being equal, the optimization will prefer slower resources with fewer miles to be paid for. This would be contrary to one of the other ISO goals, which is to avoid a formulation that makes fast resources look artificially expensive just because they can provide more mileage.⁸

4.2 The Risk of Gaming from the Inequality Constraint

The advantage of the equality formulation is that it would avoid the following game that can arise in the ISO's proposal. Consider a situation in which the requirements for capacity and mileage are set so that the resources selected to meet the regulation capacity constraint (1) could consistently provide more mileage (in terms of the sum of their capacity weighted by m_i) than is required by the mileage constraint (3). As a result, not all the mileage that could be provided by the selected regulation capacity is 'scheduled' (and therefore "costed out") by the market software. The units with the highest offer prices for mileage will only be scheduled to provide the minimum mileage specified by constraint (5). Since the floor set by constraint (5) is only somewhat more than 1/6 the actual mileage capability of a conventional resource and a much smaller proportion of the mileage capacity of a fast resource, there is a potential for resources to take advantage of that situation to strategically offer their capacity in a way that could result in high BCR payments for excess mileage costs, inflating costs to consumers and potentially harming market efficiency.

The best way to explain this is with a simple example. For instance, say that the regulation capacity required by constraint (1) is $Req_{reg} = 100$ MW, while the mileage requirement on the right side of (3) is 375 MW for an hour. Say that two resources, $i=1,2$ are selected to provide regulation capacity, and that they are scheduled to provide 75 MW and 25 MW, respectively, of regulation capacity Reg_i . Assume that unit 1 is a relatively slow resource ($m_1 = 5$) while unit 2 is

⁸ On the other hand, it is possible that only the mileage constraint will be binding, if the mileage requirement is relatively high. Then the price of mileage will also reflect the cost of capacity of the marginal source of mileage because of constraint (4), while the price of capacity itself would be zero. This outcome is also unlikely unless the CAISO substantially reduces the regulation capacity requirement so that scheduling enough capacity to meet the mileage requirement will more than meet the capacity requirement.

a relatively fast resource ($m_2 = 10$), with the former offering mileage at \$20/MW, and the latter at \$50/MW (the bid cap). The optimal mileage ‘schedule’ will then be 350 miles from resource 1 and 25 miles from resource 2 (since constraint (5) forces resource 2 to provide at least that amount). (Note that these resources could have provided up to $5*75+10*25$ MW or 625 miles, well in excess of the 375 MW that is required, so the market software chose the cheapest possible source of mileage.) The mileage price will be the cost of marginal mileage (\$20/MW).⁹ Assume that resource 1 set the capacity price at \$30/MW with its offer. As a result, if resource 2 actually produced 25 miles when dispatched by AGC, it would earn $30*25 + 20*25 = \$1250$ in revenue. If its capacity was offered at \$0/MW, its as-bid cost ($50*25 = \1250 for just mileage) would be just barely covered by its revenue. Based on the schedule from the market software, the overall estimated payments for regulation would be $30*100+20*375 = \$10,500$.

Now, consider what happens when those resources are actually dispatched by the AGC system. Assume that the overall mileage estimate of 375 MW used in constraint (3) was an accurate forecast of the AGC mileage. But because resource 2 is much faster than resource 1, if the resources were dispatched in proportion to their ramp capability over the hour, the AGC system might instruct resource 2 to provide, say, 150 miles, and resource 1 would provide 225 miles. Then resource 2 will incur a large apparent loss, as it will be paid $30*25$ for its capacity and $20*150$ for its mileage (or \$3750 total) but its as-bid cost for that capacity and mileage will be $0*25$ and $50*150$, respectively (totaling \$7500). The difference (\$3750) will be eligible for bid cost recovery. If $i=2$ is a storage device that has no other revenues from the ISO markets, then it will obtain a BCR payment in that amount. The total payment by the ISO for regulation will then be $30*100$ for capacity plus $20*375$ for mileage and \$3750 for BCR, or \$14,250. This is 35% higher than the \$10,500 regulation payments anticipated by the market software’s solution.

This opportunity for a large BCR payment arises because the use of an inequality constraint in (5) potentially distorts the ‘schedule’ of mileage from different sources. If market clearing capacity prices are significant, that potential payment provides an incentive for resources to simultaneously understate their capacity price offer and overstate their mileage price offer. If such a resource can be confident that the sum of those two offers will be less than the sum of the two clearing prices, it knows it will be selected to provide capacity and 1 MW of mileage for each MW of capacity. Its optimal bids are then to bid as low as possible to provide capacity and high as possible to provide mileage in order to maximize the amount of BCR it would receive.¹⁰ These bids are subject to that constraint on the sum of the offers, and the restrictions that the capacity offer price cannot be negative and the mileage offer price cannot exceed the ISO’s proposed cap of \$50/mile.

This opportunity to strategically distort bids to maximize BCR has two potentially deleterious effects. One is the inflation in the payments ultimately made by consumers for the provided

⁹ The mileage price is not \$50/MW, even though some mileage is acquired from resource 2 which offered that price for its mileage. The reason is that resource 2’s mileage is at its lower bound (constraint (5)). If the mileage requirement in (3) is increased, the additional mileage would come from resource 1 (@ \$20/MW), not resource 2.

¹⁰ This assumes that the unit would make money from the capacity, mileage, and BCR payments that result; if not, then it should instead bid so that it is not selected to provide capacity.

regulation, as just described. The second is that the distorted relative as-bid costs of capacity and mileage could, in turn, result in a potential distortion in choices of resources to provide regulation, possibly inflating the true economic cost of providing regulation. As an example of how this can happen, consider that there might be a third resource ($i = 3$) whose capacity bid was, say, \$35/MW and mileage bid was \$25/MW, and whose m_i was 5.5. It would not be chosen over resource $i = 2$, because its as-bid cost to supply 25 MW of regulation capacity and mileage would be $35 \times 25 + 25 \times 25 = \1500 , whereas $i = 2$'s cost was $0 \times 25 + 50 \times 25 = \1250 , but its as-bid cost of providing the 150 MW of mileage that resource 2 actually provided would have been only $\$25 \times 150$ or $\$3750$ compared to $\$7500$ for resource 2. If the schedule that minimizes social cost was actually the two slow resources, the ability of resource 2 to distort its bid in a way that undercuts resource 3 instead harms market efficiency. This shows that the potentially large distortions in estimates of mileage by the market software could result in important misallocation of resources.

It is important to notice that this game is not an exercise of market power, in the sense that a large player is taking advantage of its ability to affect prices. Very small but fast resources can play this game just as readily as large fast resources. For instance, imagine that resource 2 in the above example actually consisted of ten 2.5 MW resources rather than a single 25 MW resource. Each of the 2.5 MW resources would have precisely the same incentive to minimize its capacity bid and maximize its mileage bid in an attempt to collect excessive BCR. Moreover, these inefficient outcomes can arise from purely cost-based bidding without any gaming behavior at all. There is nothing in the example concerning resources 1 and 2 that would be different if resource 2's actual capacity costs were zero and its actual mileage costs were \$50. If capacity clearing prices are material, there is a potential for cost-based bidding to lead to inefficient outcomes and excess BCR costs. The role of possible gaming is simply to magnify the potential costs.

However, the key element that creates the potential for the inefficient bidding in the example is that the clearing price of capacity is materially different from zero. If the clearing price of capacity is zero because most of the costs currently recovered in capacity payments would be shifted in the mileage bid, the potential to exploit this discrepancy will not exist. Lots of resources underbidding their capacity costs and inflating their mileage offers would also tend to drive down the capacity price and eliminate the potential for material distortions. Given the rather fundamental change in clearing prices that will accompany the shift to pay for performance bidding it is hard to predict the likely level of capacity clearing prices and hence to assess the realistic potential for inefficient outcomes.¹¹

Therefore, it is important that the ISO closely monitor the level of capacity clearing prices, and the level of BCR payments to assess whether this kind of inefficient outcome is arising, either as

¹¹ That this gaming opportunity would persist if either capacity or mileage bids from slow resources are high can be seen as follows. If the cleared capacity bid price is high (say \$100/MW), then a fast resource would likely be able to provide a low capacity bid (say \$40/MW) and a high mileage bid (say \$50/MW) and be accepted, being 'scheduled' to provide its minimum of 1 MW of mileage per MW of capacity (by constraint (5)). Or if the cleared mileage price bid is high (say \$50/MWh), the fast resource could just undercut that and be accepted. The if the AGC obtains a much larger amount of mileage from the fast resource, that resource would be eligible for BCR.

a result of gaming behavior or from cost-based bidding behavior interacting with the discrepancy between the scheduling criteria and AGC instructions

4.3 Possible Measures to Prevent the Game or Mitigate Its Effects

We describe six alternative approaches for reducing the scope or impacts of the inefficient outcomes identified in the previous section. These are offered as possibilities to be considered should the problem described above become important in the market. With the possible exceptions of the first and third proposals, each has large significant disadvantages that mean that we do not recommend full implementation of any of them at the present time.

1. The first approach would be to lessen the scope for the game by tightening the lower bound on mileage ‘scheduled’ in constraint (5) of the Proposal. Presently, it is equivalent to one-sixth the regulation capacity, which is much closer to zero than it is to the likely typical use of a regulation resource.¹² For example, this lower bound could be raised or made proportional to m_i rather than set equal to 1 MW of mileage for each MW of capacity. We recommend that some such adjustment be considered as testing progresses of the proposed design between now and implementation.
2. A second approach would be the logical extreme of the first proposal in which the lower bound (5) could be set equal to the upper bound (4), which is the equality constraint proposal in Section 4.1. This would eliminate the potential for this kind of inefficient outcome entirely. But as pointed out in Section 4.1, it would also result in more difficult to interpret mileage prices which might also be zero more often.
3. A third approach would be to determine the m_i parameter for individual generators based upon a reasonable estimate of expected actual mileage under AGC, rather than using the ramp capability. Unlike approaches 1 and 2, this would tighten the constraints by lowering the upper bound (4) rather than raising the lower bounds (5). For instance, the ISO could multiply all m_i 's by some ratio of expected actual AGC mileage to potential mileage. Unfortunately, the MSC does not have access to past ISO mileage data, and cannot assess whether realistic m_i values would be much less than the maximum ramp capabilities that the ISO proposes to use. However, if the ISO's proposed values are biased but not too far off, it will also take market participants some time to figure out *how* they are biased. The ISO would then have some time to adjust the values as long as it tracks actual mileage and can makes adjustments promptly without a year long stakeholder process.
4. A fourth approach would be to eliminate the excess costs by not allowing BCR to be collected on the difference between ‘scheduled’ mileage from the market software and actual AGC mileage when a high mileage bid results in a generator's ‘scheduled’ mileage being at the lower bound (constraint (5)). This would eliminate gaming as there would be

¹² The one-sixth figure arises from noting that the formulation of constraint (5) assumes that it applies to a one hour time interval, whereas the capacity of a regulation resource is based on a 10 minute interval. So to obtain the ratio of the lower bound on miles per 10 minute interval to regulation capacity, that lower bound is divided by six.

no advantage to understating capacity costs and overstating mileage costs in order to exploit the difference between ‘scheduled’ and AGC mileage. A serious problem with this approach is that it would also not compensate resources who bid their costs and were moved far more by AGC than assumed in the evaluation. This outcome could drive these resources out of the market and perhaps discourage development of fast resources that would be particularly likely to be impacted by the discrepancy. Furthermore, this would set a perhaps unwelcome precedent of discriminating among different types of as-bid costs, allowing some to qualify for BCR but disallowing others.

5. The fifth alternative would be to allow BCR, but only on the difference between a defined “default mileage bid” (DMB) and the mileage price, rather than on the difference between the mileage price offer and the mileage price. If the actual AGC mileage exceeds the ‘scheduled’ mileage (say 3.5 MW vs 2 MW of miles per MW of regulation capacity), then the extra 1.5 miles would be eligible for BCR. The dollar amount eligible would be $3.5 - 2 = 1.5$ MW times the DMB minus the mileage price. This might be viewed as being similar in philosophy to local market power mitigation, in that energy bids are re-set to the default energy bid for any incremental energy provided in the all-constraints run relative to the competitive constraints-only run of the market software. However, the complexity that would be introduced by the need to define DMBs would be significant, and might not even have any impact on excess BCR costs, since even cost-based bidding could lead to excess BCR costs.
6. A final alternative would be to alter the AGC algorithm so that it dispatches regulation using a cost-based criterion rather than simply on ramp capability. If the cost of mileage was used to prioritize resources for dispatch by the AGC, then its dispatch would be more likely to align with that ‘scheduled’ by the market software, and the scope for excess BCR to arise either from gaming or cost based bidding would be lessened. However, any such change would require significant changes in the AGC software as well as careful consideration of the reliability and area control error consequences of such a change.

5. The Need for Consistency in Defining Requirements and Mileage Multipliers

There is an important interaction and a need for consistency between the definitions of the amount of mileage to be acquired by the market software (‘demand’, the right side of constraint (3) in the Proposal) and the amounts that resources can provide (‘supply’, defined by the constraints (4) and (5)). Inconsistencies can distort prices and increase the inefficient outcomes described in Section 4.2 whether arising from cost-based bidding or gaming.

There are two broad alternatives for consistent definitions of demand and supply of mileage. One is to base both on expectations (in the sense of a probability-weighted average) of mileage needs and use, so that the mileage amounts in the market software are reasonably representative of the amounts that AGC would be expected (on average) to use. The other alternative is to base the requirement on a target amount of mileage capability well in excess of the average so that there is enough capability in case an unexpectedly high amount of mileage is needed. Then the

supply should be based on the maximum amounts of mileage that resources could supply, which may be much greater than average mileage that AGC would demand of them.

An issue with the ISO's present proposal is that the definitions or requirements and supply capabilities appear inconsistent. On one hand, the mileage requirement is to be based on expected values, in particular the previous week's experience (at least until more accurate forecasting methods are developed). On the other hand, the supply that resources can provide is based on the maximum that could be supplied. In particular, estimates of m_i are to reflect certified ramp capability. This combination of a relatively low mileage requirement with a high potential supply means that the constraints (4) and (5) will be loose, in the sense that only some of the regulation resources chosen for capacity will have more mileage 'scheduled' than their minimum amount (in constraint (5)). This is likely to provide a distorted picture of the relative amount of mileage that will actually be provided from each resource. In particular, there will be significant opportunity for high mileage resources to play the game described in Section 4.2. And even if high mileage resources do not play the strategy of increasing their mileage bids in an attempt to inflate BCR payments, if their mileage costs are higher than for slower resources, such BCR payments may still be made frequently. Further, distortions in the choices between slow and fast units, as described at the end of Section 4.2, may still happen.

Another risk from this inconsistency is that faster resources may be undervalued and not scheduled to provide regulation capacity to the extent they should. This may occur because there will likely be intervals when AGC would have demanded a lot of mileage from fast resources if they were available, but perhaps only relatively slow resources were been scheduled because the expected mileage constraint of the market software was easily met by the ramp capability of the slow resources.

If the CAISO will want to have enough mileage capability to meet the mileage requirement when it is more than expected, the CAISO should set a target for procuring mileage capability that is greater than the expected mileage. (How much larger would depend on the balance between increasing the cost of resource of acquisition versus the probability and consequences of having less mileage capability than the AGC system would like.) Setting a target in this manner would also be more consistent with the use of values of m_i that reflect ramp capability rather than expected AGC usage.

For a given set of mileage price offers, this increase in the mileage requirement would increase the market clearing price of mileage (as more mileage would be 'scheduled'). However, we anticipate that offerers of regulation would recognize the fact that 'scheduled' mileage would then be more, on average, than actual AGC mileage, and so would likely adjust their mileage price offers down to reflect the probability of that AGC won't use all the ramp capability. As a result, the mileage price and payments might not be more than would be the case if expected mileage requirements were used in constraint (3) and m_i was based on some estimate of expected use by AGC of individual resources. This does require more sophistication on the part of regulation capacity owners, who must then factor in the probabilities of different levels of AGC utilization when bidding mileage.

6. Other comments

6.1 Bid Caps

In theory, adding a mileage payment (with an associated cap of \$50/MWh) to the present regulation capacity payment (with its \$250/MWh cap) will, in effect, increase the overall cap on the offers to provide regulation and thus potentially the cost of regulation. However, under present market conditions, the price of regulation very rarely approaches that cap, and so the effective raising of the cap is unlikely to affect prices.

However, if the need to manage intermittent renewable supplies results in a tightening of the regulation market, then this effective increase in the cap might make a difference. This could be a concern especially for fast resources if it turns out that they are able to effectively manipulate market outcomes in the manner described in Section 4.2. This possibility should be monitored by the ISO.

6.2 Payments for Energy Provided by Regulation

Stakeholders have variously proposed that imbalance energy for regulation be paid-as-bid if the as-bid cost exceeds the real-time energy price (as PowerEx has requested), or that it be paid the generator's unit-specific default energy bid (SCE). We believe that the present system of paying the market clearing price is less likely to have negative impacts on market efficiency than either of those two proposals.

In the case of the pay-as-bid proposal, we believe that pay-as-bid systems, in general, have poor incentives for cost-based bidding and should be used sparingly or, preferably, not at all. In this particular circumstance, since energy bids would not factor into the market software's selection of regulation capacity, there is no market discipline placed on energy bids from energy generators. This could incent generators to make energy price offers well in excess of their marginal cost, at least for the portion of their capacity that they anticipate will be taken for regulation. We have been unable to identify a reasonably uncomplicated and nonarbitrary way to include energy costs into regulation bids.

Market participants can account for projected losses in the energy market by factoring them into regulation capacity or mileage bids, which would allow the ISO to take the costs into account in scheduling regulation.

In the case of the default energy bid proposal, we believe that its adoption would discourage resources from offering into the regulation market relative to other markets. Although we understand that the regulation market is highly competitive, we see no reason to impose rules that could negatively impact participation in the regulation market.

5.3 Cost Allocation

We agree with the ISO and with several stakeholders that cost allocation for ancillary services, including regulation, should be based on causation to the extent possible. We believe that the appropriate forum for examining these issues and possibly restructure the ISO's mechanisms for recovering the costs of services is through the comprehensive review scheduled for later this year, and not in the context of a compliance filing such as this. We look forward to participating in that discussion.

6. Conclusions

FERC Order 755 has mandated a bid-based mileage payment for regulation resources. Any system that calculates such a payment must deal with the fundamental contradiction that arises when, on one hand, price is calculated by a market optimization whose resulting schedule for mileage is neither physically or financially enforced while, on the other hand, the mileage that regulation providers are paid for is based on an entirely different, non-optimization-based algorithm. Prices then are very likely to be inconsistent with actual operations. This leads to incentives for non-cost based bidding and the possibilities of significant bid cost recovery payments and inefficient scheduling. These potential problems will be present with any system responsive to FERC Order 755, short of a wholesale redesign of the automatic generation control system to dispatch regulation based on as-bid costs; such a redesign is impractical at this time and would have uncertain reliability consequences.

We are highly supportive of the general approach taken by the ISO's pay for performance proposal in which prices for regulation capacity and mileage are calculated as shadow prices to capacity and mileage requirements constraints, respectively. Compared to a system in which only one composite regulation product is acquired, whose price would then be decomposed into capacity and mileage portions, this two commodity approach offers fewer opportunities for owners of regulation to game the pay for performance system at the expense of customers and market efficiency

However, the ISO's proposal cannot resolve the fundamental contradiction just indicated, nor can any other design short of a redesign of the AGC system. Some risks of market gaming or inefficiencies are inevitable, and therefore must be anticipated and monitored. We are concerned about the possibility for excess BCR costs, arising either from cost-based bidding or gaming, because of systematic differences that are likely to arise between the 'scheduled' mileage in the market software for such resources and the actual mileage resulting from the automated generation control (AGC) system. The possibility of a wide range of possible mileage 'schedules' between the scheduled regulation capacity and the hourly ramp capability of a resource provides a potential for AGC instructions to differ substantially from the mileage quantities assumed in economically evaluating resources to schedule regulation, and for excess bid-cost recovery payments. If the regulation market remains very deep and capacity prices are low, then the impact of this discrepancy would be small. Given the distinct possibility that this will be the case, we believe that there is time to make adjustments in the pay for performance system as experience is gained.

Among the possible actions that could be taken to mitigate this gaming opportunity, we believe that adjustments to the parameters of the mileage payment mechanisms, in particular to the required amount of mileage and the ramp capability of the resources, could to some extent reduce the scope for inefficient scheduling outcomes. Presently, the lower bound for ‘scheduled mileage’ in the optimization software (equal to one-sixth of the capability associated with the accepted regulation capacity bid) together with a total mileage requirement that is likely to be low relative to regulation capability results in a relatively loose set of constraints on the market software’s mileage ‘schedule’. This has the benefit of ensuring a positive mileage price, but increases the scope for the inefficient scheduling outcomes we have described above.

Unfortunately, there is insufficient information available at the present time about the mileage requirements of the system compared to the potential mileage resources could provide to ascertain whether this scope would be large or small under the present proposal. Historical data from the ISO’s AGC system on total mileage as well as variations in mileage obtained from individual resources is needed to understand how the proposal would perform, and we understand that this data is being developed. This data should be obtained for a wide range of system conditions and carefully evaluated before implementation of pay for performance, with appropriate adjustments made in the mileage requirements and/or the procedures for calculating the resource mileage multipliers. In particular, increases in the lower bound for mileage (5) should be evaluated as part of the pre-implementation testing of the pay for performance mechanism.

If the analysis reveals that the scope for inefficient scheduling or regulation and its effects on the market are potentially large, implementation of more elaborate design changes should be considered. Hence, the ISO should closely monitor developments in this market and be poised to quickly adjust the parameters of the pay for performance mechanism if problems arise; it should not wait a year to make such adjustments if problems appear soon after implementation. It is desirable that the parameters can be adjusted quickly if problems emerge.

Regarding other issues, we offer the following conclusions:

- We believe that the requirements for mileage should reflect not average AGC mileage for the system, but higher values to accommodate occasions when significantly more mileage is needed.
- We do not believe that the effectively higher overall bid cap for regulation that results from having separate caps for mileage and capacity will yield higher prices, at least under present market conditions.
- Imbalance energy for regulation should be paid the real-time energy price because of the inefficient incentives that the pay-as-bid or default energy bid alternatives would provide owners of regulation resources.
- Finally, we look forward to addressing cost allocation issues in a comprehensive manner for all ancillary services, and do not recommend their separate consideration in this initiative.