

FINAL
Opinion on Integration: Market and Product Review, Phase 1

by

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

Members of the Market Surveillance Committee of the California ISO

December 8, 2011

1. Introduction

The Market Surveillance Committee (MSC) of the California Independent System Operator has been asked to provide an opinion on the three components of the Nov. 4, 2011 Draft Final Proposal for the Renewable Integration: Market and Product Review, Phase 1.¹ The purpose of that proposal is to identify modifications to the market rules that could be quickly implemented in order to facilitate integration of variable renewable power sources into the ISO markets. The fundamental concern being addressed is the potentially inadequate amounts in real-time of economic bids during periods of very low (i.e., negative) prices. Such inadequacies are experienced occasionally at present, and are anticipated to increase significantly in frequency as more renewable resources come on line. The economic consequences of inadequate economic bids when prices are negative is that additional within-CAISO thermal resources are required to be committed rather than less expensive imports so that decreases in loads in real-time can be accommodated by dispatching those resources downward. With a deeper pool of economic bids, including bids from renewable generators, it will become less expensive for the ISO to procure decreases in generation when needed under negative prices.

The Proposal involves three parts intended to encourage more economic bids in real-time. The first is revisions to the Participating Intermittent Resource Program (PIRP) designed to have contracting parties realize more directly the value of their real-time production so as to motivate them to bid more flexibly, in particular to offer to reduce output when the value of their output is very low and in fact negative. The second is a lowering of the bid-floor so that negative bids below the present -\$30/MWh floor can be made. Presently, that floor is smaller (in absolute value terms) than the value of supplemental sources of revenue such as production tax credits and renewable energy credits received by renewable generators. When these sources of revenue are included, renewable generation can produce positive net revenues even when the energy price is below the present bid floor of -\$30/MWh. The third part would alter Bid Cost Recovery (BCR) procedures to allow for separate calculation of BCR in the day-ahead and real-time markets (the latter defined for the purpose of BCR as including as bid costs in both the real-time dispatch and the Residual Unit Commitment market). The intent is to prevent generators from

¹ <http://www.caiso.com/Documents/DraftFinalProposal-RenewableIntegrationMarket-ProductReviewPhase1.pdf>

being motivated to refrain from offering their resource into the real-time dispatch in order to avoid either being dispatched in a manner that reduces profits earned in the day-ahead market or offsets real-time profits against day-ahead BCR payments.

Earlier versions of the PIRP, bid floor, and BCR portions of the proposal have been discussed during MSC meetings, most recently during the March 18 and September 30 2011 meetings. In addition, MSC members have participated in stakeholder calls and have reviewed stakeholder comments submitted to the ISO.

In general, the MSC strongly supports the goal of encouraging economic bids that would allow for downward generation adjustments in response to negative real-time prices. The importance and value to the system and, ultimately, consumers of such increased flexibility will grow significantly as the penetration of variable renewable rises. We support the general direction of the ISO's proposals as likely being effective in advancing that goal in the short-run, which is the focus of Phase 1 of the Market and Product Review process. However, we recommend some modifications to specific aspects of the proposals.

More specifically, we support the proposed PIRP revisions as a step in the right direction of making contracting parties bear the costs of the variability of their resources. We note that complete abolition of PIRP might result in significantly more economic bidding than the present proposal, which would maintain PIRP. However, without knowing the specific terms of existing contracts between PIRP resources and load-serving entities (LSE), nor knowing the extent to which such contracts would be modified to permit more flexible bidding, we cannot conclude that complete abolition would definitely result in appreciably more economic bids when prices are negative.

We also support the decrease in the bid floor to $-\$150/\text{MWh}$. However, we recommend that the proposed further decrease to $-\$300/\text{MWh}$ not be automatically implemented, but rather be recommended only if it is concluded that insufficient economic bids were elicited by the $-\$150/\text{MWh}$ floor; that appreciably more economic bids would result from dropping the floor further; and that negative side effects of having dropped the bid floor to $-\$150/\text{MWh}$ are absent or acceptably small.

Finally, we support the separate calculation of BCR for the day-ahead and real-time markets, and anticipate that it will remove important disincentives to bidding in real-time. However, we are unable at this time to conclude with confidence that the Performance Measure and Persistent Uninstructed Energy Check features of the proposal will function as intended. Although we understand and appreciate the intent of those features, there is neither sufficient detail regarding the parameter values that would be used in applying these features, nor sufficient testing data to allow us to reach a conclusion about their functioning. It is important to ascertain that those features are (1) effective in discouraging strategic behavior aimed at increasing BCR payments, (2) while not inadvertently yielding large decreases in BCR payments for normal deviations from dispatch instructions. Such decreases would undermine the goal of encouraging more resources to participate in the real-time dispatch. It is possible that significant changes to the basic features as proposed would be necessary to accomplish these goals. Therefore, the procedures should be subjected to extensive testing and possible refinement. We understand that the

ISO will be conducting such tests, and we look forward to participating in the review of those results and their implications for the design of these features.

Since the issuance of the draft opinion, the ISO's proposal for the December Board meeting has been revised. The revised proposal is consistent with our recommendations on the bid floor and BCR, while revisions to PIRP have been deferred.

2. Participating Intermittent Resource Program (PIRP)

The fundamental features of PIRP are the use of ISO forecasts to determine the real-time output schedules of PIRP resources, paying PIRP resources for their real-time output based on these schedules rather than their actual real time output, and the socialization of the cost of the differences between those schedules and the actual real-time output of PIRP resources. Real-time markets in ISO systems act as balancing markets, where the ISO dispatches flexible resources based on their bids and offers to ensure that supply and demand remain constantly balanced. Individual generators or loads have committed to production schedules in advance, and if they deviate by either producing more or less than those schedules, the ISO balances the system by adjusting the output of other suppliers. A central principle of the balancing market process is that the costs (or benefits) of such deviations from schedules are paid (or earned) by the parties responsible for the deviations. These costs reflected in the real-time price: the marginal cost of adjusting the system dispatch at the moment of the deviation.

Conventional resources have both more predictability and, at in many cases, more flexibility than do most renewable sources of supply. In part because of this, the CAISO has treated some renewable resources differently than other resources in the balancing market process. Instead of being paid the value of their real-time output on an interval-by-interval basis, PIRP resources are paid the real-time price for their hourly schedule, with an adjustment for the net difference between their actual and scheduled output over the month. Thus, a wind supplier that ended up producing 1 more MWh during a month than projected in aggregate for the month would be paid for that incremental megawatt based on the average of the real-time prices over that month. This is the case even if that producer was short of production by 99 MWh in one hour in which prices were very high and had excess production of 100 MWh in another hour during which prices were very low.

If prices in all periods were the same, or close to the same, this averaging would not be a significant issue. The problem is, not all intervals are alike. In some intervals the system may badly need more power and in others may have too much. These values are captured in the interval prices. In this context the PIRP program creates several incentive problems. First, if a renewable source subject to PIRP is allowed to schedule its power, it would have an incentive to schedule more during high price periods and less during low-price periods. It would earn the higher prices on its schedules but only pay the average of all prices for the deviations from its schedules.² For this and other reasons, the CAISO produces the schedule for PIRP resources, rather

² For example, consider a PIRP generator that actually produced zero output in real time. If this generator scheduled 1 MW of output in what it expected to be the highest price interval of the month, it would then earn the scheduled price for that 1 MW. Since it does not generate at all, its monthly imbalance is -1 MW, which would be charged at the average of all monthly prices. From this, the generator nets the difference between the highest and

than leave this to the resources themselves. Second, there is a distortion in incentives to take actions to depart from the forecast-based schedule. An interval price of \$500/MWh is signaling that the system highly values increased output, yet a PIRP generator would not be paid this price for its incremental output. On the other hand, a large negative price would signal the opposite, but the PIRP generator would not have an incentive to decrease its output.

It is worth tabulating the net value of these deviations for purposes of comparison. Say that the average monthly real-time price is \$70/MWh. If a unit produces 1 MWh excess during an interval with a negative price (e.g., -\$30/MWh), and is 2 MWh short in an hour with a high price (e.g., \$100/MWh), the costs to the system of the difference between actual and scheduled output are reflected in the sum of the deviations times the respective price difference of the prices, $+1 \text{ MWh} * -\$30/\text{MWh} - 2\text{MWh} * \$100/\text{MWh} = -\$230$. However the generator would pay for its net monthly imbalance of 1 MWh net at the average monthly price of \$70/MWh. The portion of the \$230 cost to the system that is not paid by the generator ($\$160 = \$230 - \$70$) represents a subsidy to PIRP generators. Note that it could work the other way, as for example if a unit could be 2 MWh short when the price is -\$30/MWh while overproducing by 1 MWh when the price is \$100/MWh. The generator would have, again, a net negative deviation of 1 MWh for the month, and so it would pay the ISO $\$70/\text{MWh} * 1 \text{ MWh} = \70 . Meanwhile, the cost to the system would actually decline by $-2 \text{ MWh} * -\$30/\text{MWh} + 1 \text{ MWh} * \$100/\text{MWh}$, or \$160. That is, the generator pays the ISO \$70, even though system costs are actually \$160 lower because of the deviations. However, deviations from schedule tend to be negatively correlated with real-time prices in the case of PIRP resources because the deviations of different renewable facilities will likely be positively correlated. As a result, when renewable production is lower than projected, real-time prices are more likely to be higher, and conversely extremely low prices are more likely when renewable production is unexpectedly high.

We recognize that the PIRP program provides some value to generators. For wind producers whose output is completely random, with no control over the output, the volatility of real-time prices represents a tangible source of risk. When averaged over, say, a month, we believe that this uncertainty is likely to be considerably less than risks associated with the overall variability of wind output and real-time prices.³ Nonetheless, for uncontrollable wind sources, PIRP can be viewed as a form of insurance against the cost of uninstructed deviations. However, in the terminology of the insurance analogy, the premiums are not actuarially fair. This is much closer to crop-insurance (a Federally subsidized instrument) than it is to auto insurance. Additionally, a common issue with insurance is that it can create incentive problems known as “moral hazard.” If a firm is insulated from the financial risks and consequences of its decisions, they will

average real-time price. Symmetrically, if a generator would produce 1 MW in a low price interval, it would prefer to schedule zero MW, allowing it to get paid the average monthly price for the resulting 1 MW upward deviation.

³ We anticipate that the month-to-month total imbalance costs do not greatly fluctuate, since the hour-to-hour imbalances will cancel out to a large extent. As a result, the ISO’s PIRP proposal represents a reallocation not so much of risks, but of the average of the hourly imbalances over a month, which is a much more stable quantity. We have not seen evidence that fluctuations in these costs are anywhere near as important for a typical wind generator as, for instance, fluctuations in revenues due to variations in monthly average wind production or in the ISO’s market prices.

be less inclined to take steps to avoid or mitigate those consequences. In this context, the moral hazard is the lack of incentive to adjust output in response to changes in real-time prices when a failure to adjust output would be most costly to the system. There may be ways to subsidize the cost of the insurance while still preserving the marginal incentive to respond to current interval prices.

The most important feature of the PIRP portion of the Nov. 4 Draft Final Proposal proposes the following change to the current PIRP program: that for individual PIRP resources (under both existing and future contracts), uplift costs associated with PIRP will be allocated to a scheduling coordinator (SC) of a LSE designated by the resource rather than shared by all load.⁴ This proposal has evolved considerably over the last year. The ISO originally proposed to eliminate the PIRP program for new generation but faced strong opposition. The ISO then proposed to allow LSEs and generators to keep signing new contracts under the old PIRP rules until FERC approved the change. Finally, what the ISO proposes to do now is to effectively eliminate the subsidy elements of PIRP by allocating all of the deviation costs to the SC/LSE that contracts with the PIRP resource (unlike today where they all load shares the cost equally).

We believe that this assignment of costs internalizes the imbalance costs arising from variability/forecast errors to the wind generator and the SC/LSE it has designated, and is much more desirable than the current systemwide sharing of the costs. In theory, this internalization then provides incentives to parties who can do something about the imbalance costs. The parties can renegotiate contracts (if necessary) to provide incentives to bid more flexibly in the face of negative prices, and perhaps, in the long run, to build new facilities whose power outputs can be better forecast. In the case of new PIRP contracts, the SC/LSE and resources will be incented to negotiate terms that will result in more efficient operation.

However, some stakeholders and DMM are concerned that this reallocation of imbalance costs will not incent more efficient operation in the case of existing PIRP contracts. In that case, the contracting parties (the resource and its SC/LSE) would be leaving money on the table by failing to facilitate bidding in a way that allows the resource to be turned off when the negative real-time power price falls below the negative of the production tax credit and value of renewable energy credits. Possible reasons for this include the inflexibility of existing contracts and the difficulties involved in renegotiating them, as well as the absence of appropriate communications and control equipment in older wind turbines.

On the other hand, in the case of new PIRP contracts, we do not see why allocating the costs to the SC/LSE would not create the right contracting incentives going forward. If the renewables market is operating reasonably competitively, the buyers of PIRP generated power benefit from the “insurance premium” subsidy as much as the sellers of PIRP power. There is good reason to expect that if PIRP were ended completely, the resulting increase in costs would be reflected in

⁴ This might normally be the LSE that has a purchased power agreement (PPA) with the resource, but this is not necessarily the case under the proposed PIRP revisions. Presumably, if another LSE takes on the cost allocation of the uplifts, it would have to be compensated by the PIRP resource for providing this financial service.

higher contract prices.⁵ Either way, the SC/LSE would be the likely party to bear the brunt of the costs (explicit or implicit) of PIRP.

The ISO cannot prevent the SC/LSE from simply choosing to absorb these costs, but under the CAISO's present proposal they would no longer be able to assign them in part to other SC/LSEs. We anticipate that this will incent SC/LSEs to negotiate purchased power agreements with individual resources that provide for submission of economic bids when prices are sufficiently negative. With the costs of imbalances internalized in contracts, the contractual terms should reflect the division of those imbalances. For instance, this could result in higher PPA payments from the SC/LSE if SC/LSEs transfer imbalance costs back to the generators who might be in a better position to manage them (for instance by reducing output when prices are very negative). But if the generator would rather not manage those costs, it can arrange to pay for the equivalent of insurance from third parties or, more likely, accept lower PPA prices in exchange for the counterparty SC/LSE choosing to pay those costs instead. This can be negotiated by the parties, and the results of those negotiations will not be distorted by the present policy of sharing the imbalance costs across all load.

Note that there are two somewhat separate issues here. The first is the distribution of the costs of the imbalance insurance provided by PIRP. The second is the degree to which that insurance leads to moral hazard problems by muting incentives to respond to changing system needs. Even if it is decided that renewable firms require subsidies for insuring against the risks of imbalances (which we doubt, as these risks are likely to be less than other market risks), this can be done in a manner that still preserves incentives to respond to system prices, which will improve market efficiency.

We conclude that the ISO proposal to reallocate PIRP costs to the SC/LSE has promising potential to mitigate the distortions in bidding decisions and perhaps even investment decisions that arise from the current practice of sharing the costs of uplift charges across all load. A strong, but perhaps insufficient incentive will be provided to SC/LSEs and resources involved in present PIRP contracts to adjust bidding procedures and perhaps the contracts themselves to permit more flexibility in responding to negative real-time prices. Unfortunately, we lack sufficiently detailed information regarding the structure of typical existing contracts, either in terms of payments or control of operating decisions, to make an informed judgment as to whether complete elimination of PIRP would lead to a better outcome than the ISO's current proposal to shift the cost allocation. But that we believe that either the ISO's proposal to change the cost allocation or complete elimination of PIRP would provide needed improvements relative to the status quo. We note that more megawatts of new wind resources are expected to come on line under new PIRP contracts than already operate under existing PIRP contracts, so it is most important to get incentives right for the new resources.

⁵ In general, the incidence of a new cost or tax is borne by the less price responsive (or elastic) part of the market. In this case, the "demand" for renewable power can be considered quite inelastic due to the requirements of California's renewable portfolio standard (RPS). As long as the RPS is enforced, it is likely that the costs of this PIRP adjustment will be ultimately be passed on to SC/LSEs whether they are first assigned to resources or assigned to SC/LSEs directly. The important difference to the current method is that SC/LSEs will no longer be able to spread the charges of their own PIRP related costs to other SC/LSEs.

3. Bid Floor

Recall that all the measures contained in the ISO's proposal are motivated by a goal of creating better incentives for firms to be more flexible and responsive in both their bids and their operations. The adjustment of the bid floor is another measure aimed at furthering this goal. There is a periodic and growing problem with "overgen" conditions where the system has too much energy offered at zero or even slightly negative prices. In other words, too many units are unwilling to reduce their output even though they are paying, through a negative price, for the privilege of continuing to produce power.

This rigidity comes from several sources, including a desire to maintain fixed schedules for institutional or contractual reasons. However, various policies directed at promoting renewable energy are playing an increasingly prominent role. An unfortunate side-effect of these policies is that they reward production of renewable MWh no matter whether they help or harm system performance. These policies include production tax credits (PTC) and renewable energy credits (RECs) that make production of renewable energy profitable even when the energy price is negative. Firms can pay to inject energy into the transmission grid when energy prices are negative but earn even more back in the form of PTC and REC revenues. During times of large negative prices, obliging the power system to accept renewable energy will increase system operating costs and can increase emissions relative to a better designed renewable subsidy program.⁶ If the objectives of renewable energy policy is to lower costs and emissions, and to decrease renewable energy costs through learning curve effects, incentives for production that are blind to its effects on the system work against the first two objectives, and do nothing for the third. We would support revisions to renewable support schemes that would avoid the perverse operational impacts of those policies.⁷ These schemes could be designed to maintain the financial levels of support for renewables while reducing or eliminating these counter-productive side-effects.

Given the renewable policies that are in place, we accept the need to lower the bid floor to encourage flexible response by generation during overgen situations. Given the present magnitude of tax credits and the cap on the price of renewable credits in California, -\$150/MWh should be low enough to elicit economic bids from renewable generators. There should be very few circumstances where outside payments and other costs are large enough for a unit to operate profitably at -\$150.

However markets, particularly balancing markets, do not always operate as seamlessly and smoothly as one would hope. The CAISO balancing market has often suffered from a lack of flexible bids from conventional units and renewable units, not to mention load. There appear to be contractual or perhaps institutional preferences that result in resistance to adjusting units

⁶ B.F. Hobbs, *Transmission Planning and Pricing: Lessons from Elsewhere*, Transmission Policies to Unlock America's Renewable Energy Resources, PESD, Stanford University, Sept. 15 2011; G. Oggioni, F.H. Murphy, Y. Smeers, "A Stochastic Version of Market Coupling with Wind Penetration," International Conference on Operations Research, August 30-September 2, 2011, Zurich, Switzerland.

⁷ An example of such a reform would be to modify California law so that the renewable generation would not count towards the 33% requirement during dispatch intervals when the local price is negative. If the output in these intervals did not count towards a renewable portfolio standard (RPS) requirement, LSEs would not be interested in buying the renewable energy credits associated with this output. The incentive to generate when the price is negative would be significantly reduced, although federal production tax credits would remain.

from their schedules. Further, with some contracts, the firm in control of operating the facility may not be the one exposed to the negative energy prices. All these factors could work to limit the level of response bids that the lower cap may produce. At the same time, the lower floor will expose market participants to more price volatility.

Because we anticipate that circumstances will be rare when a unit can operate profitably at -\$150/MWh, we also recommend that further reductions in the bid floor below -\$150/MWh *not* be automatic, as it is in the CAISO proposal, which would reduce the floor to -\$300/MWh after one year. We believe that additional reductions should occur only after study of the impacts of the initial drop to -\$150/MWh (the “opt in” approach) to determine if it was effective in eliciting more flexibility, and if there were no significant unanticipated negative effects. This includes addressing questions such as the following:

- What increase in the quantity of dec bids occurred, especially at times when they were most needed? That is, will the ISO have received enough decremental offers at prices above -\$150 so that it is able to balance the system without curtailments or undue use of regulation?
- Which resources did those additional economic bids come from, and were they are reasonably cost-reflective? This information might shed light on contractual and other barriers to submission of economic bids.
- Was there unanticipated inefficient behavior that occurred when prices fall towards -\$150?

This recommendation of ours has also been made by the Department of Market Monitoring (DMM), as well as several stakeholders in their comments, and we agree that some caution is justified.

One gaming concern that has been raised by DMM is that generators in generation pockets could use lower bid floors to extract larger profits if real-time transmission deratings require that the generators be dispatched down below their day-ahead schedules in real-time. Although transmission outages that reduce the transfer capability out of generation pockets between the day-ahead and real-time markets are unlikely to occur often, the ISO’s local market power mitigation (LMPM) procedures would not protect the market from generators with large amounts of local market power from exploiting this situation. There is no market efficiency reason to allow such generators to extract large rents in these circumstances. If indeed such behavior is judged by DMM to be a serious risk, we would support the intent of a proposal for a targeted mitigation rule to be applied specifically in the context of such real-time transmission deratings that reduce transfer capability out of specified concentrated generation pockets below the level of day-ahead schedules. But such a rule should not use this possibility to constrain the dec offers of the all intermittent resources around the state.

4. Bid-Cost Recovery (BCR)

4.1. General Comments

In general, the economic rationale for bid-cost recovery is that in the presence of non-convex costs,⁸ marginal costs based on the last accepted supply or demand bid in the system (or possibly a penalty for a violated constraint) may not “support” the cost-minimizing solution. If prices support a solution, this means that individual market participants cannot profitably deviate from the solution’s schedule. However, the presence of non-convexities such as start-up and minimum load costs or prohibited operating zones mean that a generator who is scheduled in the least cost solution might lose money under the solution’s prices. If prices do not support the least-cost solution, then market parties would be incented to deviate from ISO schedules or to self-schedule in order to yield higher profits for themselves. The result, however, of such deviations would be higher overall costs for the market.

Several theoretical schemes have been proposed to define supporting prices when non-convex cost functions are present in power markets. These schemes can result in something very similar to the BCR rule of paying a lump sum to a generator if its gross margin is negative to erase the incentive to shut down rather than to produce.⁹ A further implication of theory is that separable decisions should, in theory, earn separate BCR-type payments. As one example, once day-ahead commitments are in place, a generator considering whether to offer in the real-time market should not face the possibility of the cost-minimizing real-time schedule resulting in a loss. Therefore, BCR payments for real-time should not depend, in theory, upon BCR payments made for day-ahead schedules. As another example, a short-start unit might be scheduled to start and shut-down for both the morning and afternoon ramps. Since these are separate decisions, an incentive is, in theory, needed to ensure that the generator is willing to start-up in both times, implying separate BCR payments for each on-off cycle.

The ISO’s proposed changes to the BCR procedures would provide for separate calculation of BCR day-ahead and in real-time, as opposed to the present system of paying based on overall gross margin summed over all the markets. This is consistent with the above principle of separate BCR calculation for separable decisions. However, the proposal would retain the daily netting of costs and revenues across the 24 hours of the market. The direct purpose of the proposed change is to eliminate the elements of the current CAISO design that cause generators to self-schedule their day-ahead market schedules in real-time so that they cannot be dispatched down. One incentive for this behavior is that at present, bid cost recovery is calculated over day-ahead and real-time, so that if a generator has a profitable day-ahead schedule and then is

⁸ A convex total cost function is either linear or upward bending function of output (i.e., nonnegative second derivative). Nonconvex cost functions either bend the other way in some regions, or have discrete jumps (as in start-up costs) or prohibited regions.

⁹ For instance, R.P. O’Neill, P.M. Sotkiewicz, B.F. Hobbs, M.H. Rothkopf, and W.R. Stewart, Jr. “Efficient Market-Clearing Prices in Markets with Nonconvexities,” *Euro. J. Operational Research*, 164(1), July 1, 2005, 269-285; W.W. Hogan and B.R. Ring, “On Minimum-Uplift Pricing for Electricity Markets,” March 19, 2003; P.R. Gribik, W.W. Hogan, S.L. Pope, *Market-Clearing Electricity Prices and Energy Uplift*, HEPG, December 31, 2007

dispatched uneconomically in real-time, it makes less money than if it were not on dispatch. No other ISO/RTO has such a rule because they wanted to avoid precisely this kind of outcome.

In particular, a generator that earns a positive gross margin (revenues minus variable costs) in the day-ahead market might be reluctant to provide adjustment bids in the real-time market because such bids might be accepted but result in losses in the latter market. For instance, this could happen if a short-start unit was scheduled for the morning ramp in the day-ahead market, and then in the real-time market would be scheduled (perhaps at a loss) during the afternoon ramp. The loss in the afternoon period would not yield a BCR payment if the day-ahead morning schedule was sufficiently profitable. In the reverse situation, in which there is a loss in the day-ahead market (revenues less than costs in the absence of bid-cost recovery), the incentive to bid in real-time is blunted by the effect that at least some of the positive gross margin that could be earned in real-time would be taken away by decreasing the BCR that compensates for the day-ahead loss. This would occur because BCR is presently based on the sum of the negative day-ahead margin and the positive margin in real-time.

Therefore, the proposed separate calculation of day-ahead and real-time payments is likely to incent more adjustment bids in the real-time market. This would improve market efficiency, in terms of reducing the as-bid cost of meeting load. This does not necessarily mean reductions in consumer costs because BCR payments might rise. Simulations reported Section 4.6 in the ISO's proposal of the effect of the BCR rule change indicate that an increase in BCR payments on the order of 20% could occur; however, that analysis necessarily assumed no change in bidding behavior. In assessing the consumer impact of these improvements in market efficiency, it is important to recognize that if the Phase 1 changes are successful in eliciting more economic bids for reducing output when prices are negative, then BCR payments could actually decrease, or at least not increase as much as those simulations would indicate. This is because those additional decremental bids would reduce the magnitude of downward price spikes and hence reduce BCR payments relative to what they would have been otherwise.

Besides the market efficiency benefits of removing disincentives for suppliers to submit offer offers to be dispatched down in real-time, there are at least two other possible advantages of the proposed change. One is that the various schemes that were being used to extract uplift costs from the CAISO and were addressed by the March and June 2011 filings perhaps took advantage of the present system which bases BCR on the combined day-ahead/real-time gross margin. This opportunity would be lessened by the proposal. Another possible advantage is that the current pattern of rising minimum load offer costs may represent an effort by generators to reduce the amount of day-ahead rents that can be expropriated by the current BCR rule by raising their as-bid costs to be closer to the market price. Although this may or may not actually be the case, if it is true, then this is another inefficiency induced by the current rule that would be reduced or eliminated by the proposed change.

4.2. Performance Metric and Persistent Uninstructed Energy Check

The BCR portion of the Phase 1 proposal would apply a performance metric that would scale certain components of the bid cost recovery calculation based on deviations from ISO dispatch instructions. This prorating process is intended to remove the incentive, for instance, for a generating unit to receive a day-ahead BCR payment based on a day-ahead commitment, and then to declare an outage so that the unit would not actually run but would still receive the payment. However, the performance metric only considers uninstructed deviations within a single real-time interval. Therefore, the ISO proposes to augment the performance metric with a real-time calculation of a persistent uninstructed energy index that would disqualify real-time energy from the real-time BCR calculations in the case where generators choose to deviate consistently over several periods, yielding a greater deviation between actual operation and system cost-minimizing dispatch than can result from just one interval's deviation. The check would construct a 'counterfactual' or hypothetical series of operating levels that would have occurred if the generator had adhered to the operators' instructions.

We generally support the need for the proposed performance metric and persistent deviations check, and the general philosophy behind their calculation and application. They are likely to be more effective than the present BCR procedures in avoiding potential BCR payment inflation from intentional deviations.

However, these changes represent significant departures from the previous BCR procedures at the ISO, and indeed at any other RTO or ISO. For this reason, it is important that the procedures, as well as the particular parameter values to be used to implement them, be subject to careful testing to ensure that they will work as intended. In particular, will they effectively guard against intentional inflation of BCR payments arising from unscheduled deviations? And, at the same time, will they avoid penalizing innocent behavior by prorating BCR payments in response to normal scheduling inaccuracies or errors in a way that would undermine the goal of encouraging more resources to participate in the real-time dispatch? In the absence of reporting of thorough testing of the procedures, we do not have enough information at this time to state with confidence that the procedures will solve the problems they are intended to solve without negative effects on other aspects of the market.

We recognize that there are benefits to implementing all these Phase 1 changes at once, since separation of day-ahead and real-time BCR together with decreases in the bid floor might increase gaming opportunities. However, it is not yet clear that the performance metric and Persistent Uninstructed Energy Check will be effective in addressing the potential problems relating to BCR for uninstructed output, while avoiding undesirable side effects. There needs to be evidence that these proposed procedures will in fact address the potential problems before the California ISO commits to adopting them. Therefore, we recommend that a final commitment to this portion of the BCR proposal not be made until after careful testing has been undertaken to refine the general concepts and to identify appropriate ranges of parameters, and these results have been reviewed by stakeholders. It is our understanding that the ISO will be undertaking such tests, and we look forward to participating in the review of those results and their implications for the design of these procedures.