



## Can we prevent the gaming of ramp constraints?

Shmuel S. Oren<sup>a,\*</sup>, Andrew M. Ross<sup>b</sup>

<sup>a</sup>Industrial Engineering and Operations Research Department, UC Berkeley, Berkeley, CA 94720, USA

<sup>b</sup>Industrial and Systems Engineering Department, Lehigh University, Bethlehem, PA 18015, USA

Available online 20 July 2004

### Abstract

Some electric power markets allow bidders to specify constraints on ramp rates for increasing or decreasing power production. We show in a small example that a bidder could use an overly restrictive constraint to increase profits, and explore the cause by visualizing the feasible region from the linear program corresponding to the power auction. We propose three penalty approaches to discourage bidders from such a tactic: two based on the duality theory of linear programming (LP) and the other based on social cost differences caused by ramp constraints. We evaluate the approaches using a simplified scaled model of the California power system, with actual 2001 California demand data.

© 2004 Elsevier B.V. All rights reserved.

*Keywords:* Auction design; Ramp constraints; Power generation dispatch

Many restructured electricity systems rely on self-commitment of generation resources rather than on central unit commitment. This structure avoids some of the incentive-compatibility problems associated with more centralized systems such as the original UK system (prior to NETA), PJM, NYPP, New England pools which involve multi-dimensional auctions allowing bidders to specify technical constraints on the dispatch. Similarly, FERC's proposed Standard Market Design allows ramp constraints to be specified. Such auctions are often susceptible to

manipulation, allowing bidders the opportunity to profit by specifying deceiving technical constraints. Unfortunately, in systems that rely on self-commitment and clear the hourly day ahead market without consideration of intertemporal constraints on dispatch, mismatches between the ISO schedule and the capabilities of generators must be made up in the real-time balancing market. Not only is this an expensive solution, it shifts a perhaps unnecessary volume of energy transactions to the real time balancing market. Furthermore, although some generation technologies hinder efficient scheduling due to their ramp constraints, and others assist with their rapid ramping capabilities, the rapid-ramping plants are not rewarded for the flexibility they bring to the system. In this paper, we explore ways to allow

\* Corresponding author. Tel.: +1 510 642 1836; fax: +1 510 642 1403.

*E-mail address:* [oren@ieor.berkeley.edu](mailto:oren@ieor.berkeley.edu) (S.S. Oren).

bidders to specify ramp rate constraints while mitigating to some extent the possibilities for deceptive bidding. We are concerned with two effects of such bidding: an increase in overall cost, and inequity of outcomes.

We deal here with a day-ahead energy market without network considerations, and allow bidders to specify ramp rate constraints as part of their offers. The decision to turn a unit on is still left to the generator who will need to absorb the startup cost and can ensure it minimum generation level using zero or negative offer prices for that amount. In order to accommodate ramp constraints, the market operator would need to clear the markets for all the 24 hours simultaneously using an optimization algorithm. This problem, fortunately, is far less complicated than a unit commitment problem since it does not introduce discrete variables into the problem. If the objective is to minimize the social cost of the dispatch, then the market clearing problem with ramp constraints can be formulated as a linear programming (LP) problem and solved by standard algorithms and software. The solution gets more complex if the objective is to minimize total procurement cost, since that criterion introduces non-convexities. It should be noted that the day-ahead electricity auction in the Spanish system allows bidders to specify ramp constraints as well as a floor on their total 24-hour revenue, which are incorporated into the market clearing formulation and solved using a heuristic algorithm [2].

One simple strategy for designing the auction to avoid manipulation is to allow ramp-rate constraints to change only once a month, or some other suitable time interval, since generator technologies do not change very rapidly. While this has little theoretical backing, its practical effect would hopefully be to prevent generators from rapidly responding to market conditions with false ramp rates, leaving them with little choice but to report their true ramp rate

Table 1  
Problem data

	Off-peak	Peak
Demand	1 GW	3 GW
Generator A offers	1 GW, \$10/MW h	1 GW, \$10/MW h
Generator B offers	2 GW, \$15/MW h	2 GW, \$15/MW h
Generator C offers	2 GW, \$25/MW h	2 GW, \$25/MW h

Table 2  
Auction results without ramp constraints

	Off-peak	Peak
Clearing price	\$10/MW h	\$15/MW h
Generator A	1 GW	1 GW
Generator B	0	2 GW
Generator C	0	0

constraints. However, it would also allow a company to lock in a misleading constraint and profit from it for a whole month. Another simple idea is to have a regulatory agency certify ramp rates, as the California regulations stated; however, this leads to only an upper bound on the rates. Bidders might be bidding into various markets, and so would need to split their ramp rates among them, giving them a valid reason to specify ramp rates lower than their certified values. The same reasoning applies to the change-only-once-a-month proposal.

The problem data given in Table 1 help demonstrate the incentive compatibility problem that can arise when ramp rates can be specified and are accounted for in the dispatch but the spot market prices for energy are based on snap shot market clearing that does not account for intertemporal constraints affecting the dispatch. For illustrative purposes we will assume that each time period has a duration of 12 h. Table 2 summarizes the auction results (for minimizing social cost) in the absence of ramp constraints.

Suppose now that generator B specifies an intertemporal constraint requiring that his dispatch level at peak is no greater than his dispatch level at off-peak. Then, the auction results under optimal social-cost dispatch are as given in Table 3. We note that the intertemporal constraint stipulated by generator B caused an increase in social cost due to displacement of cheap energy with more expensive energy, while the resulting increase in market clearing prices increased the net profits of both generators A and B. While one

Table 3  
Auction results with generator B ramp constraint

	Off-peak	Peak
Clearing price	\$15/MW h	\$25/MW h
Generator A	0 GW	1 GW
Generator B	1	1 GW
Generator C	0	1 GW

Table 4  
Financial summary

	Unconstrained	Ramp-constrained
Generator A profit	\$60 000	\$180 000
Generator B profit	0	\$120 000
Generator C profit	0	0
Social cost	\$600 000	\$780 000
Acquisition cost	\$660 000	\$1 080 000

could not fault generator A for enjoying the windfall, it is clearly inappropriate for generator B to reap extra profits by stipulating a constraint that impedes efficiency. Such a profit opportunity could motivate generators to misrepresent their ramping capability in order to drive up prices. Some market designs (e.g., the old UK system) attempt to prevent misrepresentation of constraints by barring a constrained generator from setting the clearing price. Our example demonstrates, however, that such a restriction still does not solve the problem since the constrained generator may force a more expensive unit into the dispatch and benefit from the higher clearing price set by that unit. The Spanish market design eliminates such perverse incentives by forcing generators to bear the dispatch consequences of their ramp constraints, which in our example would amount to forcing generator B out. This rule solves the incentives problem but, unfortunately, it may also unnecessarily increase the social cost of the dispatch (Table 4).

1. Visualization

We wish to visualize the feasible region and optimal solution for our small example, to learn how generator B can profit by specifying a ramp constraint. Our small example has six variables (one for each generator in each period) and two equality constraints (one for demand in each period). Using the equality constraints, we can eliminate two variables. Furthermore, generator A is always assigned 1 GW during peak regardless of us including or excluding generator B’s ramp constraint, so we can treat generator A during peak as a fixed quantity rather than as a variable. This reduces us to three variables, which are more easily visualized than our original 6. Fig. 1 shows the feasible region without generator B’s ramp constraint; it is a triangular prism and the optimal solution is marked




Fig. 1. Feasible region without ramp constraints.

with an extra circle. We start imposing the ramp constraint gradually in Fig. 2, where the increase from off-peak to peak must be under 1.5 GW for generator B. We see that the new constraint plane has cut off two corners of the non-ramp-constrained region, pushing the optimal solution so generator B gets more business during off-peak, and generator A gets less. Also, we must now take generator C’s higher price. Fig. 3 shows the feasible region (an irregular tetrahedron) once the full ramp constraint that B specified has been imposed; generator A has been pushed out of the off-peak solution entirely. Overall, we see the problem with allowing bidders to specify ramp constraints: it allows them to specify slanted constraint planes (as opposed to upper-bound constraints on power output, which are orthogonal to the axes and cannot push the solution away from other bidders).




Fig. 2. Feasible region with generator B increase limited to 1.5 GW.




Fig. 3. Feasible region with full generator B ramp constraint.

When the feasible region is viewed this way (especially as in Fig. 2), it is natural to think of the effects of the ramp-rate constraint in terms of LP sensitivity theory. We will explore this idea further in the next section.

## 2. Penalty system

If we are to allow bidders to specify ramp constraints, we should ensure that our final dispatch satisfies all of the constraints, since we do not know which are true and which are misleading. That is, in trying to make our auction more incentive-compatible, we have little leeway in the dispatch. We do have some room to adjust the payments, though. One option, explored in Ref. [3], is the Vickrey-Clarke-Groves (VCG) auction. A VCG auction effectively pays companies to bid truthfully. It pays each offer the generation cost based on the revealed parameter and an additional premium that reflect the contribution of the offer to social welfare (i.e., the difference in the optimal value of the objective function with and without the offer). While the VCG auction is incentive compatible and efficient, it may lead to revenue insufficiency and is considered undesirable due to its radical departure from the uniform price philosophy underlying commodity markets. Instead, we propose to capture the essence of the VCG approach by starting with the usual uniform (market-clearing) payments as the benchmark and impose financial penalties on companies whose ramp constraints are

active at the optimal solution. This would tend to reduce the acquisition costs, instead of increasing them as the VCG auction does. The following are desirable features of penalty systems:

- (1) Avoid under-penalizing:
  - reduce the incentive to specify misleading constraints
  - recover the increase in social cost from ramp constraints
- (2) Avoid over-penalizing:
  - more than the corresponding profit increase
  - more than the corresponding social cost increase
  - so much that a bidder's costs are not recovered
- (3) Quickly computable
- (4) Transparent
- (5) Unaffected by multiple optimal dispatch solutions

To these ends, we offer three penalty system proposals:

PP1: a penalty based on LP sensitivity theory (dual variables and right-hand-side (RHS) ranges).

PP2: a penalty based on the difference between the ramp-constrained social cost and the social cost without a bidder's ramp constraints.

PP3: a penalty based more loosely on LP sensitivity theory (dual variables, but a larger RHS change).

These systems are focused on the problem of the increased social cost; they do not directly address the potential profit increase, though we will investigate their effects on this.

### 2.1. Penalty proposal 1

The first penalty proposal (PP1) calculates penalties as follows: the full auction LP is solved, and dual variables for each active ramp rate constraint are calculated the usual way. Also, it is simple to extract for each ramp constraint a range in which the RHS can vary without changing the optimal LP basis. If a particular active ramp constraint has dual variable  $\lambda$  and allowable RHS change  $\delta$ , we propose that a

financial penalty of  $\lambda \cdot \delta$  be imposed on the company that specified the ramp constraint. In our example above, the dual variable for generator B’s ramp constraint is \$120/MW, and the allowable RHS range is 1 GW, for a penalty of \$120 000. This reduced B’s income from \$480 000 to \$360 000, exactly its value in the solution without the ramp constraint. In this case, generator B has no incentive to specify a misleading ramp constraint, unless it is in league with generator A (who benefits from B’s ramp constraint and escapes any penalty payments).

In Fig. 4, we use the viewpoint of parametric linear programming to show the three proposed penalty systems; for now, we will focus only on the first system.

The figure shows how the optimal social cost would decrease as a generator’s maximum ramp rate (in just one hour of the day) is increased. Suppose the current ramp rate is at the point marked *a*—the corresponding optimal social cost is marked with a dot. The slope of the cost curve is the dual variable  $\lambda$  for the ramp constraint. If the ramp rate increased to point *b*, the optimal basis for the LP would change, and hence the cost curve changes slope. The distance  $\delta=(b-a)$  corresponds to the allowable RHS change. Thus, the PP1 penalty  $\lambda \cdot \delta$  corresponds to the vertical distance shown by arrow 1.

Unfortunately, it is not hard to find examples where the penalty does not exactly compensate for the shifted profits. In a few cases, the penalty is too much; this tends to happen when the ramp constraint would be violated in only one period, but after adding the constraint for all periods (as it natural) there are two periods where it is binding, so penalties are charged for both periods. In other cases, the penalty is

not enough, so that a company still profits by giving a misleading ramp constraint. This can happen when the ramp constraint chops off too many corners from the feasible region. That is, our penalty system is based on the idea of Fig. 2, where the ramp-constrained optimum is adjacent to the optimum without the ramp constraints. It is this adjacency that determines how large the RHS ranges are. It is a matter of coincidence in the costs that this example works well even with the full ramp constraint, where the ramp-constrained optimum is not adjacent to the original optimum. To avoid this situation, we consider next dropping each bidder’s ramp constraints in turn.

### 2.2. Penalty proposal 2

In our second penalty proposal (PP2), we run the auction with all ramp constraints, and save that solution. For each bidder with active ramp constraints, we re-optimize having relaxed its ramp constraints. The difference in social cost between the two solutions is the penalty to the bidder. We use the social cost difference, rather than the acquisition cost difference, because social cost (as revealed by the offers) reflects the true cost of constraints and furthermore, the PP1 penalties are also based on social cost. Also, since our optimization minimizes the social cost, alternate optimal solutions can cause larger differences in the acquisition cost than bidders should be penalized for. This approach requires as many optimizations as there are bidders with active ramp constraints. However, each optimization starts with an easy known feasible point (the optimal solution with all ramp constraints imposed). Thus, the calculation of penalties is not as computationally intensive as premium computation in a VCG auction, where each new optimization (with an offer removed) starts without a known feasible point.

Referring back to Fig. 4, PP2 focuses on the drop in social cost (vertical axis) from the ramp-constrained optimum point denoted by the dot above point *a* to the non-ramp-constrained optimum cost shown on the cost curve above point *c*. If the ramp rate increased to point *c*, the constraint would be in effect relaxed, and so further increases would not change the optimal social cost, so the cost curve becomes horizontal. Since the constraint is effectively relaxed, we get the same answer as if we were to drop the constraint,

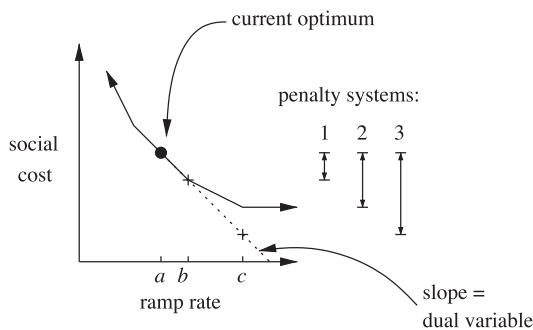


Fig. 4. The cost decreases as the ramp rate increases.

which is the heart of the PP2 system. The penalty would then be the vertical distance shown by arrow 2.

By dropping ramp constraints and re-optimizing, our penalties are not restricted by other adjacent vertices of the feasible region, so in most cases penalties will be at least as large as the penalties from PP1 (the exceptions being when PP1 penalties over-penalize because of multiple active constraints). This makes the auction more incentive-compatible. However, it is possible that a penalty might be so large that the penalized generator will end up with a deficit. Indeed, the penalty might be even greater than the income itself. In either of these cases, we can assume that the penalized company would want to withdraw its offer. We would then exclude that company and re-optimize (the question remains whether to exclude all such companies simultaneously or one at a time starting with the worst-off). Such exclusion, however, will drive up social and acquisition costs since eliminating an offer amounts to adding a constraint on the optimization. This could happen with the PP1 penalties as well, but not as often, since PP1 penalties are typically smaller than PP2 penalties. In either case, there is a trade-off between reducing acquisition cost by imposing penalties, and increasing cost by forcing out offers through excessive penalties. Another problem might occur if an offer that is forced out by a large penalty is needed for reliability reasons. If this happens on a continuing basis then a Reliability-Must-Run (RMR) contract could be enacted, but it would be harder to deal with if it happened only occasionally.

### 2.3. Penalty proposal 3

Our third proposed penalty (PP3) is again based on duality theory, just like PP1. However, it is more drastic: instead of multiplying the dual variable  $\lambda$  by the allowable change in the RHS before the basis changes, we multiply it by the RHS change needed to relax the ramp constraint completely. That is, if the dispatch without ramp constraints had a maximum dispatched ramp rate of  $c$  (referring to 4), and the ramp-constrained dispatch has a ramp rate of  $a$ , we multiply  $\lambda$  by  $(c-a)$  to get the penalty for this ramp constraint. On the figure, we see that this corresponds to the vertical distance between the original optimum point and the cross-hatch marked

on the dotted line. Due to the convexity of the cost curve, this point is lower than the optimal social cost with the constraint relaxed. This vertical distance is shown by arrow 3.

We calculate the value of  $c$  by looking at the ramp rates implied for each hour transition by the non-ramp-constrained dispatch. This means that PP3 requires two optimizations (one ramp-constrained, the other not) whereas PP1 requires only one. However, this is still faster than PP2, which requires one optimization for each bidder. A potential downside is that PP3 is more vulnerable than PP1 or PP2 to multiple optimal dispatches, since the size of the penalty depends on  $c$ , which might vary from one optimal solution to another.

An interesting note is that the value of  $c-a$  is not always positive. This can happen when the unconstrained dispatch had a large one-period jump that is then spread out by the ramp constraints to cover two or more periods. We could either ignore the penalty for that ramp constraint in those cases (in effect, let the penalty be  $\lambda \cdot \max(0, c-a)$ ), or continue using the formula  $\lambda(c-a)$ . We have chosen to continue with the simpler formula, since in one case the resulting total penalty was equal to the PP2 penalty. This is an interesting result that might deserve further study, but must be left for another time.

To summarize our proposed penalty systems, we expect (from theory and from Fig. 4) that the penalty sizes will usually fall into the order  $PP1 \leq PP2 \leq PP3$ , where the equality cases are a common occurrence. However, we have mentioned some toy cases where PP1 is larger than expected, or where a penalty is so large that a bidder would want to back out. Next, we explore the effectiveness of these penalties on a simulation of the California power system, along with the number of penalty anomalies (inverted ordering or back-out cases).

## 3. California 2001 CalECo Simulation

We tested our suggested penalty approaches using generator data from the CalECo system of Refs. [4,5], which represents a scaled abstraction of the California power system developed for the purpose of evaluating production simulation models. It includes several generators from each fundamental type (nuclear, coal,

gas, etc.), and four price-quantity energy blocks for each generator. Further details of the CalECo system are provided in the cited references and will be omitted here.

The demand data was collected from the California ISO web site by Matt Schneider. Due to difficulties in the data recording, only 330 days had full data for all 24 hours. Table 5 shows how the 35 omitted days were distributed with respect to the 31 days when load-shedding was implemented. We see that  $31/365=8.5\%$  of the total days had load-shedding and  $25/330=7.6\%$  of the days in our data set had load-shedding, so we are considering a fairly representative sample of 2001.

Because the maximum output of the CalECo generators (12500 MW total) is not on the same scale as the current California demand data, we scaled the demand data to create three data sets in which the scaled annual peak load (originally 41244 MW, on August 7 at 4 p.m.) left reserves of 1.5%, 5% and 15%. To isolate the effects of the ramp constraint mechanisms, we assumed that no bidders specified ramp constraints except for the “gas3” bidder, who specified a ramp constraint of either 150 or 200 MW/h (2.5 or 3.33 MW/min), which is clearly more restrictive than a real gas-fueled generator would face. This constraint was imposed for every hour of every day.

We used the AMPL mathematical programming language and the CPLEX linear program solver to run the auctions. For each day in each data set, we used the following procedure:

- (1) Optimize social cost without any ramp constraints.
- (2) Optimize social cost with all ramp constraints.
- (3) Compute penalties for all three proposed systems.

In all cases, the nuclear, QF and coal bidders were run at full capacity in every period, so ramp constraints for them would have been irrelevant.

Table 5  
Demand data summary

Days	In data set	Omitted
No load shedding	305	29
Load shedding	25	6

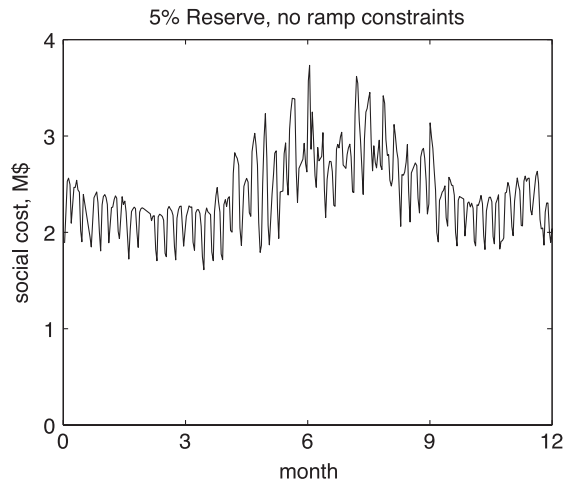


Fig. 5. Daily social cost for one demand scenario.

### 3.1. Effects of ramp constraints

Before we evaluate the penalty systems, we will explore the behavior of the simulated system without the penalties. To get a feeling for the data set, Fig. 5 shows the social cost throughout the year for one of the three demand levels, with no ramp constraints imposed (the other two demand levels look almost the same, except for scaling).

Notice that, due to the scaled demand, the financial figures are much smaller than one would expect in reality. For this reason, we will focus our evaluation

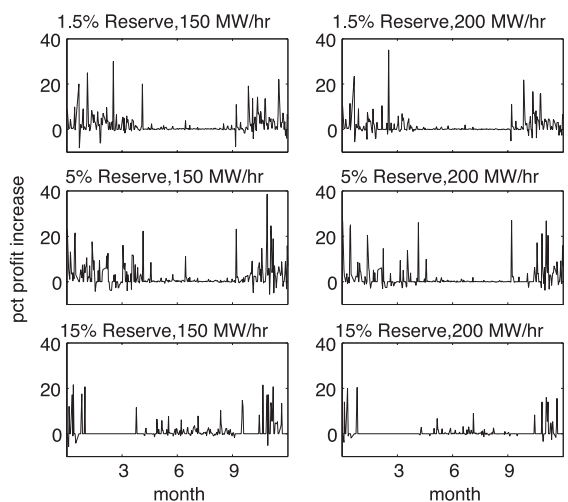


Fig. 6. Daily profit increase (in %) from no-constraint to constrained.

Table 6  
Percent of days with increased profits

Reserve (%)	150 MW/h (%)	200 MW/h (%)
1.5	75	53
5	63	38
15	28	16

on percentage rather than absolute changes. In Fig. 6, we show the percent that profit for “gas3” increased when that bidder specified a ramp constraint. This figure is broken out into categories by demand level and ramp constraint value (150 or 200 MW/h). The higher percentage increases tended to be on days without much original profit, though.

Tables 6 and 7 summarize Fig. 6, showing the percent of days with a profit increase and the average percentage increases for these days (weighted by profit amounts). We focus only on the days with an increase because a bidder that uses this scheme would try to carefully choose when to specify the constraint, so as to avoid days on which specifying a ramp constraint would result in reduced profit.

3.2. Effects of penalty systems

Up to this point, we have described the effects of misleading ramp constraints. Now we turn to the effects of the proposed penalty schemes. Fig. 7 shows, in percentages, the net profit increases (after PP1 penalties) for the various scenarios, plotted against the gross increases. Fig. 8 similarly shows the results for PP2 and Fig. 9 for PP3. Ideally, these graphs would be scattered along a horizontal line through zero (the net profit increase would be negligible). Unfortunately, we see that this is not the case. However, we do see that the penalties do fall in the expected order, that is, PP3 pulls more net profit increases toward zero than PP2 does, which in turn is better than PP1. Interestingly, there are days where the gross increase is negative, and then a further penalty is applied. This is acceptable to the extent that one of

Table 7  
Percent profit increase (when positive)

Reserve (%)	150 MW/h (%)	200 MW/h (%)
1.5	0.45	0.45
5	0.66	0.71
15	0.76	0.67

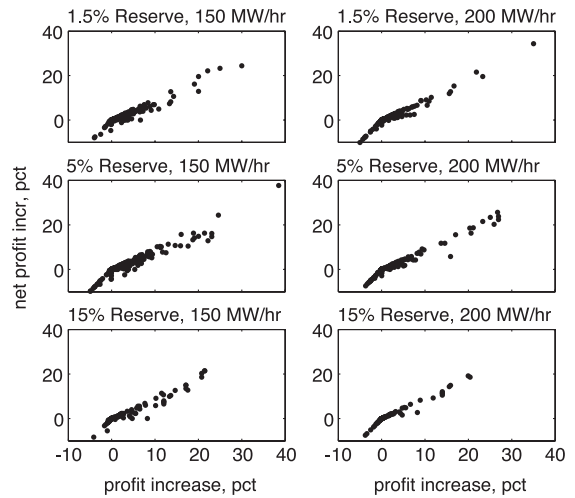


Fig. 7. Net (after-PP1-penalty) vs. gross percentage increase in profits.

our aims is to recover increases in social cost due to ramp constraints.

Fig. 10 shows, for PP1 penalties, the percentage of the increase in social cost that the penalties recovered, as the year goes on. We would hope that the values would concentrate around 100%. Fig. 11 shows the same results but for PP3 penalties. It has been cropped at the top of the vertical scale—a few days managed to recover over 10 times the increase in social cost and the champion was 21 times the




Fig. 8. Net (after-PP2-penalty) vs. gross percentage increase in profits.



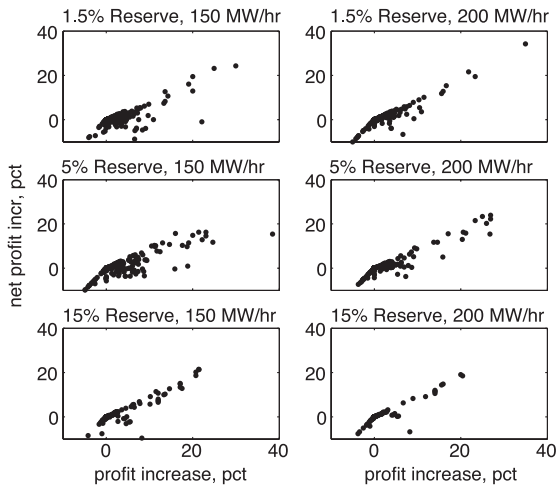


Fig. 9. Net (after-PP3-penalty) vs. gross percentage increase in profits.

increase. For PP2 penalties, the amount recovered is always exactly 100%, when only one bidder specifies a ramp constraint. It is not clear that this would be true when more than one bidder has an active ramp constraint; the amount recovered might be less than or greater than the total social cost difference.

Finally, we note that our concerns about over-penalizing (so much that a bidder would prefer to back out) were not an issue in our simulated data set—it did not happen in any of our six scenarios (ramp rates and demand levels) as it did in a few toy




Fig. 10. Percent of social cost increase that PP1 recovers.




Fig. 11. Percent of social cost increase that PP3 recovers.

examples. Also, our expectation that  $PP1 \leq PP2 \leq PP3$  was true in most cases. It was always true that  $PP2 \leq PP3$  and Table 8 summarizes how many days had anomalous results for PP1.

We see more anomalous results when reserves are tight and (in all but one case) more anomalies when the ramp constraints are more restrictive.

#### 4. Variations on Spain’s system

Spain’s electricity market rules favor a heuristic solution procedure, rather than a process based on mathematical programming. The market rules for Spain’s system include the following restriction [1] (p. 28):

In any case, when the owner of a production unit which includes the rising/start-up or descending/stop load gradient condition in an electric power sale offer, the market operator shall assign the producer a lower quantity of power than the latter would have received if it had not included the cited condition.

Table 8  
Days when PP1 was larger than PP2 or PP3

Reserve (%)	PP1>PP2		PP1>PP3	
	150 MW/h	200 MW/h	150 MW/h	200 MW/h
1.5	29	20	7	8
5	23	6	3	2
15	4	3	1	1

This prevents a company from gaining business by specifying ramp constraints, which is our aim, without specifying penalty payments. Inspired by this rule, we consider some variations: when a ramp constraint is specified, the constrained dispatch must be compared to the unconstrained dispatch in one of these four ways:

- (1) The power quantity cannot increase in any hour, or
- (2) The sum of power dispatches over the day (i.e., energy dispatch) cannot increase, or
- (3) The income cannot increase in any hour, or
- (4) The total income for the day cannot increase.

The first two are fairly easy to implement as simple linear constraints once the initial dispatch (without ramp constraints) is obtained. The first becomes a set of 24 constraints and the second becomes a single constraint. The third and fourth variations are much harder to implement, since the income in any period is the product of the market-clearing price and the dispatch quantity, and so is a nonlinear term. Furthermore, we have to add binary variables to the LP formulation to calculate the market-clearing price in this context. It is possible to eliminate the nonlinearity by noting that the market-clearing price must come from the set of offers, and creating a constraint for each combination of possibilities, but this becomes unwieldy very quickly. Overall, from the market perspective (ignoring implementation difficulties) it seems that restrictions on the total daily income make more sense than hourly incomes, due to cost differences between hours of the day. Also, income constraints seem better than MW allocation constraints, since the bottom line is profit rather than power generation (although in Spain power generation may have an indirect effect on profit due to stranded cost payments.)

While any of these four approaches sound fair, there are two other predicaments that should be considered: they can make the problem infeasible and they depend heavily on the initial solution. Electric power dispatch problems are notorious for having multiple optimal solutions. If the solution chosen as the initial one gives a particular company only a small allocation (of power, energy or income), while another solution gives it a larger share, it seems unfair to restrict that company to the smaller of the

two. To avoid this problem, though, we might have to optimize once for each company, trying to give it as big an allocation as possible while maintaining (near-) optimality. This would significantly increase the computational requirements of the auction process.

## 5. Conclusions

The PP1 penalties have the advantage of being easily computed from a single run of the auction LP, along with having the usual economic interpretations of dual variables. However, we have seen that PP1 is not particularly effective in recovering social cost losses from generators due to strategic specification of ramp constraints. The cost-based (PP2) penalties require many optimizations to be run, but they perform better on social cost recovery. Neither PP1 nor PP2 performs very well at removing the incentive for any particular bidder to specify misleading ramp constraints. The variations proposed for the Spanish approach are all subject to the problem of multiple optimal solutions, and the two that make the most sense financially are the most difficult to implement. Further research is needed into the problem of incentive-compatible auctions with ramp constraints. However, at this point, it seems that in order to prevent gaming, generators should be restricted in how often they can restate such constraints.

## Acknowledgements

This work was supported by the Power Systems Engineering Research Center (PSerc) and by the Electric Power Research Institute.

## References

- [1] Compañía Operadora del Mercado Español de Electricidad, S.A. Electricity market activity rules, April 2001. Available: <http://www.omel.es/en/normativa/mreglasconadhesion.htm>.
- [2] J. Contreras, O. Candiles, J.I. de la Fuente, T. Gomez, Auction design in day-ahead electricity markets (republished), *IEEE Transactions on Power Systems* 16 (3) (2001 Aug.) 409–417.
- [3] B.F. Hobbs, M.H. Rothkopf, L.C. Hyde, R.P. O'Neill, Evaluation of a truthful revelation auction in the context of energy markets with nonconcave benefits, *Journal of Regulatory Economics* 18 (1) (2000 July) 5–32.

- [4] R.B. Johnson, S.S. Oren, A.J. Svoboda, Equity and efficiency of unit commitment in competitive electricity markets, *Utilities Policy* 6 (1) (1997) 9–20.
- [5] C. Marnay, T. Strauss, Effectiveness of antithetic sampling and stratified sampling in Monte Carlo chronological production cost modeling, *IEEE Transactions on Power Systems* 6 (2) (1991 May) 669–675.



Shmuel S. Oren is Professor of Industrial Engineering and Operations Research at the University of California at Berkeley and a former chairman of that department. He is the Berkeley Site Director of PSerc, a multi-university PowerSystems Engineering Research Center sponsored by the National Science Foundation and industry members. His academic research and consulting focus on planning scheduling and economic analysis of power systems, and in

particular on issues concerning market design for competitive electricity. He published numerous papers in this area and has been a consultant on electricity restructuring issues to numerous public and private organizations in the US and abroad. Dr. Oren holds a PhD and MS in Engineering Economic Systems from Stanford and a BSc and MSc in Mechanical Engineering from the Technion, Israel. He is a Fellow of the IEEE.



Andrew M. Ross received the BS degree from Harvey Mudd College, CA, USA in 1996 in Mathematics and the MS and PhD degrees from the University of California, Berkeley, CA, USA in 1997 in Operations Research. He is currently an Assistant Professor at Lehigh University. His research focuses on non-homogeneous queueing systems that are common in the telecommunications industry.