3 March 1997

Report to the California Trust for Power Industry Restructuring

Activity Rules for the Power Exchange

Phase 3: Experimental Testing*

Market Design Inc.
Prepared by Robert Wilson, Vice President

Executive Summary

The scope of the work stipulated for Phase 3 during February includes experimental testing of the activity rules developed in Phase 2, as well as identification of problems and proposed remedies. The experimental program includes construction of a laboratory prototype by February 14, followed by a series of tests and demonstrations by February 28. This program was undertaken by Professor Charles Plott, who is a prominent expert on experimental studies of markets. He directed the construction of the prototype by H.Y. Lee, and he designed and conducted the tests at the Caltech Laboratory for Experimental Economics and Political Science. The prototype was completed by February 14, and experiments were conducted over the following two weeks, often with several trading sessions per day. The design followed current practice in studies of market mechanisms. The subjects were Caltech students whose entire remuneration consisted of their trading profits. Professor Plott conducted demonstrations on February 21 for members of the PX Team, and on February 28 for members of the TAC. He is also submitting a companion report with additional detail.

The first task was to establish whether the PX Protocol can be implemented in a working prototype. The second task was to establish whether the auction's iterative process converges, the rate of convergence, and the character of its dynamics. The third task was to measure the efficiency of the auction outcomes. Each task was divided into studies of single and multiple markets, and cases without and with fixed costs. The multiple markets correspond to the PX's 24 hourly markets for next-day delivery; the fixed costs correspond to the start-up and no-load costs incurred by thermal generators. Additional topics included the role of withdrawals, substitution among markets (as in the case of hydro supplies), and sensitivity to parameter specifications (such as the minimum bid decrement). Some tests were conducted using demand patterns and supply portfolios representative of the California mix prepared by London Economics using data from the CEC and FERC.

The main conclusions from these studies are the following:

- **Implementability.** We had no difficulty implementing the PX Protocol. The software requirements are straightforward. Subjects in the experiments had no difficulty understanding and following the procedural rules of the auctions.

- **Convergence.** In all tests the auction converged. All subjects tried to game the system but these strategies proved ineffective. After some experience, several subjects concluded that simply bidding their costs is optimal, which accelerates convergence. We conclude that the activity rules succeed in suppressing gaming behavior or rendering it ineffective.

- **Efficiency.** In most tests the auction ended with an outcome that was within a few percent of perfect efficiency. The final clearing prices and quantities were close to the theoretical equilibrium prices, even with few bidders. The exceptions were that inefficiencies occurred in tests that included either a supplier with significant market power or one with supplies that could be allocated costlessly among the markets. We conclude that activity rules cannot supplant measures to mitigate market power.

- **Rate of Convergence.** Progress is substantial in the first five or six iterations, residual inefficiency is small after eight to twelve, and full convergence often occurs in ten to twenty. Nevertheless, in extreme cases, as when the bid decrement is small, convergence can require forty iterations. Because the PX might restrict the number of iterations to as few as twelve over two hours, measures are required to accelerate convergence or to terminate the auction after progress has slowed sufficiently or when time expires; or, the allowed time might be increased or the software altered to enable more iterations or continuous bidding. We conclude that there are sufficient measures available to close the auction without significant inefficiencies.

- **Fixed Costs.** The tests with fixed costs that must be recovered from multiple markets showed comparable efficiency. Subjects learn quickly to stay active in those markets where prices are sufficient to recover their fixed costs. The dynamics follow the scenario predicted by London Economics: subjects initially load their fixed costs into their bids in each market, but then later prorate them among the markets in which they remain active. With this strategy, withdrawals are minimized and inefficiencies due to premature withdrawals are rare.

London Economics' companion report addresses these and additional topics. In particular, they conclude that inclusion of additional constraints on operational feasibility increase the number of iterations required for full convergence.

Our summary conclusion is that the PX Protocol is a viable design for an energy market, and the efficiency of its outcomes is impressive. The numbers of bidders and markets in the tests were small, and we did not replicate the daily repetition of the market, but we found no fundamental impediment to full-scale implementation. Further work in Phase 4 should refine the design to accelerate convergence and assure a timely close.

1. Review of the Activity Rules

In the absence of activity rules the auction outcome could be inefficient. Bidders could wait until the final iteration to offer serious bids, which prevents early price discovery and thereby prevents bidders from identifying their optimal hours of operation – which is essential due to the start-up and no-load costs of thermal generators. The purpose of activity rules, therefore, is to encourage early serious offers so that price discovery proceeds steadily throughout the iterative process. Their design is subject to the restriction that they cannot impair efficiency; in particular, they cannot constrain suppliers who choose to offer their actual costs.

The "standard" activity rules used for the experimental tests are summarized in Appendix A. They are based on the principle of revealed preference. The Exclusion Rule is key: a bidder cannot offer later a price that improves a previous clearing price that was not improved at the first opportunity; i.e., in the next iteration. Thus, if a supplier declines to improve the previous iteration's clearing price then we infer that this price is below the supplier's cost for that increment of supply, so the supplier is precluded from offering a lower price later. This rule is complemented by four additional routine procedural rules, stated here in the form applicable to suppliers (the rules for demanders are analogs). The Opening Rule requires that all available capacity is offered in the first iteration. The Revision Rule restricts revised prices to those less than the previous clearing price by at least a specified decrement. The Withdrawal and Closing Rules require that withdrawals are irrevocable, and they preclude withdrawals after the final iteration.

The effect of these rules on a supplier is to require an irreversible decision. If its offered price in the previous iteration exceeded the clearing price, then in the current iteration it must offer a price less than that previous clearing price or forego all later opportunities to do so. If its cost is sufficiently low then the supplier's best strategy is to revise its offered price; otherwise, it's better to decline, in which case it cannot later revise its price unless the clearing price rises higher. If all suppliers offer their actual costs then the auction ends after the second iteration, since no offers are revised.

When suppliers bid strategically by offering prices above their actual costs, several iterations are required to drive their revised offers down to their costs. The resulting competitive process involves only those suppliers near the margin. The extra-marginal suppliers must revise their offers (or be frozen out); when they do so they become infra-marginal, thereby making some previously infra-marginal suppliers extra-marginal, and now these too must revise their offers.

The rate of convergence is driven by the difference between the current clearing price and the equilibrium price. When this difference is large (relative to the elasticity), one side of the market is "long" by a large amount. If it is the supply side then extra-marginal suppliers rejected (or rationed by the Rationing Rule) and they must revise their offers or be frozen. When these suppliers revise their offers, they eject a large number of previously infra-marginal suppliers from the merit order. This ensures a decrease in the clearing price by at least the amount of the specified price decrement. The dynamics of this process are elaborated in Professor Plott's companion report.

When the difference is small, however, the imbalance may be small too and the price need not change for one or more iterations (the PX Team that visited the Lab called this "stuttering"). The difference can be small either because the clearing price is close to the equilibrium price, or because the supply elasticity is small. In either case the increase in the total gains from trade from further iterations is small, and zero if demand is inelastic too. In California, the supply elasticity is expected to be large (except perhaps in peak hours in the summer) and the demand elasticity is small. Consequently, one expects rapid progress in the first few iterations, after which the trading gains diminish rapidly, indicating that it may suffice to terminate the auction by invoking a convergence criterion.

Example: Suppose that demand is inelastic at 1000 MWh and at the current clearing price of 20 $/MWh aggregate supply is flat over the range 900 to 1200 MWh. Then 200 MWh is rationed and the suppliers offering this amount must in the next iteration offer 19 $/MWh or less, due to the specified decrement of 1 $/MWh. If they all elect to do so then in the merit order their revised offers displace those who previously offered prices in the interval between 19 and 20 $/MWh, including the unrationed 100 MW at 20 $/MWh. If the supply offered at prices less than 19 $/MWh exceeds 800 MWh then the new clearing price is 19 $/MWh, and otherwise the new clearing price is between 19 and 20 $/MWh – and surely less than 20 $/MWh if the supply previously offered at prices below 20 $/MWh exceeds 800 MWh.

This sort of example occurs when the initial 20 $/MWh clearing price is substantially more than the equilibrium price. At the equilibrium price the imbalance is nil and no rejection or rationing occurs, but if the equilibrium price is only slightly less than the current clearing price then the amount rejected or rationed is small and the amount displaced in the merit order by revised offers is also small, so there is a greater prospect that the clearing price changes by less than the price decrement. This feature indicates that for practical purposes it may suffice to adopt a criterion for terminating the auction when the percentage of rejected or rationed offers is sufficiently small.

Convergence can be accelerated by specifying a large price decrement for revised offers. This has the immediate effect that clearing prices move in large jumps between iterations. It also has a strategic effect that further accelerates convergence. A supplier realizes that if its offer is above its cost by less than the magnitude of the decrement then it will be unable to revise its offer; consequently, there is a stronger incentive to offer revised prices close to its cost.

The Exclusion Rule is stated above for the case that the auction proceeds in discrete iterations. If the auction is accelerated by allowing continuous bidding then this rule can be stated in terms of the time interval allowed for submission of a revised offer. The February 21 report provides further elaboration on this and other variants of the basic activity rules. One should be cautious about these variants, however, since in the two weeks allowed for the experiments it was possible to test only the basic set of rules.

2. The Experimental Program

The experimental program was designed to address several issues. A preliminary step was to verify that the PX Protocol is implementable. This was accomplished by constructing a prototype sufficient for several bidders (e.g., 12) and several parallel markets (e.g., 2, 3, or 4), although in principle the software is capable of much larger numbers (one test was run with 20 markets). Further, the initial tests verified that subjects easily understood the rules and regularly submitted offers that conformed to these rules. No fundamental impediments to a full-scale implementation were identified, and indeed the software requirements are straightforward. Bids were submitted manually at keyboards so iterations were slow, but since the eventual implementation will allow computerized submissions this delay can be avoided. The one lacunae found in the PX Protocol was an inadequate specification of how to handle cases without a unique clearing price, as can happen when a supplier specifies a minimum load. This deficiency was patched by using the lowest price for which supply is not less than demand.

The next step was to establish how the Protocol works in a single market, with and without fixed costs, and with and without demand-side bidding. The first week of testing verified that fixed costs did not impair the efficiency of an isolated market. Demand-side bidding produced no complications: clearing prices did not converge monotonically in this case, but non-monotonicity did not disrupt convergence. A significant conclusion from this series of tests was that the competitive process and the factors affecting the convergence rate match the theoretical predictions. In particular, the subjects' strategic bidding slowed convergence but did not prevent near-efficiency at the close. Some tests were run with data representative of the California mix, prepared by London Economics using data from the CEC and FERC.

The second week addressed three central issues posed by the peculiar features of the PX. These were run with multiple markets representing demand configurations corresponding to peak, offpeak, and shoulder periods; also, to focus on the issues, demand-side bidding was mostly excluded. The first issue was the effect of market power. As expected, the activity rules did not mitigate market power: a large supplier can sustain a clearing price above the equilibrium price by withholding supply, and at the margin it can capture the difference between its cost and the next higher cost in the merit order. The second issue was the effect of the Opening Rule on supplies that can be freely allocated among multiple markets: the tests showed that efficiency can be impaired by the restriction that capacity cannot be reallocated among the markets in later iterations. This test was imperfect because subjects were apparently unaware that they might initially offer amounts in the several markets that exceeded their total capacity and then later withdraw from (or freeze their offers in) those markets with low prices, which is a strategy that overcomes the restrictions imposed by the Opening Rule. We concluded that subsequent work in Phase 4 might consider an activity rule for total-energy portfolios; one candidate is described in the February 21 report, but based on the successful record of the predominately hydro NordPool system this need not be the first priority.

The most important tests in the second week studied the role of start-up costs when there are multiple markets. The issue was whether convergence would be "top down" and therefore efficient with few or no withdrawals, or "bottom up" and therefore potentially inefficient if some suppliers withdraw prematurely when faced with low initial prices when in fact they could operate profitably at the final prices. This issue is essentially a behavorial question. One hypothesis is that a supplier's offer in each market will include its entire fixed cost in the first iteration, and thereafter this fixed cost will be apportioned among the markets in which the supplier remains active – this strategy implies that withdrawals are largely unnecessary because a supplier can exit a market by freezing its offers above the clearing price. The second hypothesis is that a supplier will offer prices on an incremental-cost basis in each market, hoping that the eventual clearing prices will be sufficient to recover its fixed cost, and withdrawing otherwise – this strategy implies that withdrawals are important, and that efficiency depends on the order in which suppliers elect to withdraw. The simulation studies by London Economics showed that, due to the flatness of California's aggregate supply function based on incremental costs, the second hypothesis implies a potentially significant inefficiency if withdrawals are premature due to myopic expectations about the final clearing prices.

The results from the tests conducted to examine this issue support the first hypothesis. With no prompting, the subjects invariably followed the top-down strategy. Consequently, the selection of suppliers, and the markets in which they remained active, were accomplished by freezing offers, and there were no inefficient withdrawals in the test runs. Thus, the tests produced no evidence that there might be an inherent tendency for incremental-cost offers that could cause inefficiencies due to premature irrevocable withdrawals. This conclusion is reinforced by other observations that conform to the predictions from the simulations conducted by London Economics; e.g., offpeak prices converge first and quickly and shoulder prices last – in some cases requiring many iterations to settle down.

3. Conclusion

In Phase 2 we were asked to "fill in the blanks" in the PX Protocol by suggesting activity rules that would suppress gaming and promote price discovery during the iterative process of the auction. In Phase 3 we were asked to assess the reliability of the design principles in predicting actual outcomes in experimental tests. The set of activity rules in Appendix A was the candidate studied in the experimental program during the February.

The test results from over forty sessions indicate that with this amendment the PX Protocol can be implemented in a small-scale prototype, and that typically it converges to an outcome that is within a few percent of perfect efficiency. Although the number of iterations required for complete convergence could exceed the limits imposed in the 1/1/98 implementation, there are ample tools to accelerate convergence (e.g., a large price decrement and/or continuous bidding) or to conclude the auction after progress has slowed (e.g., a convergence criterion based on a measure of residual inefficiency). In any case the auction outcome is potentially feasible after every iteration, and the welfare losses from premature termination are likely to be small. A crucial test was whether subjects would adopt the top-down strategy to handle fixed start-up costs, thereby preventing inefficiencies from premature withdrawals, which indeed they did.

The efficiency obtained in practice will include aspects not considered in the experimental design. As London Economics has emphasized, the correct measure of efficiency is based on suppliers' opportunity costs, not running costs. Because the PX is a forward market, intertemporal efficiency requires taking account of expectations about subsequent prices in the inc/dec, ancillary services, hour-ahead and real-time balancing markets. This is a major reason why the one-part schedules used in the PX Protocol are sufficient: little is added, and perhaps some subtracted, by multi-part bidding.

The potential inefficiencies seem to be these four: market power (if it is not mitigated), premature withdrawals (if some suppliers adopt the incremental-cost strategy), premature termination (if necessary), and inefficient allocation of total-energy portfolios among the hourly markets (due to the restrictions of the Opening Rule). Each of these can be handled by appropriate measures developed in Phase 4. In particular, we recommend that market-power mitigation not interfere with the remarkable efficiency of the auction. The monitoring of market power should ensure that the irrevocability of withdrawals is not used by large suppliers to stalk smaller suppliers at the margin, as described in the February 21 report. The termination rule should be based on an estimate of the welfare effect to ensure that losses are insignificant, and experimental tests should measure the effect of end-game strategies.

With these provisos, we conclude that the PX Protocol is a viable market design. As far as we can tell from the prototype, the auction is capable of high efficiency and little prospect of successful gaming. No firm conclusion can be drawn about the full-scale implementation in 1998, but the present evidence is that the theory and the design principles accurately predicted the outcomes of the tests undertaken.

Appendix A

Standard Activity Rules Used for the Tests

The following "standard" version of the activity rules is the one used for the experimental tests. No attempt was made to test the several variants described in the February 21 report. The Rationing Rule was "first come, first served" based on the time stamp of each new or revised tender. This version is stated for supply tenders; symmetric rules apply to demand tenders. The tenders are assumed to be schedules that are step functions; modifications are required for piecewise-linear schedules.

Tenders: Each step of each tender is a binding offer to trade at any price not less than the offered price. Each tender remains in force until it is withdrawn or validly revised by the trader, or rejected by the PX. A revised tender replaces the previous tender for the same portfolio. At the close of the auction, those steps with prices above the final clearing price are rejected; ties at the clearing price are resolved via the Rationing Rule. The remaining steps are accepted, and each becomes automatically a binding contract, with the PX as the counter-party, for the tendered or rationed quantity at the final clearing price – except a step at the margin, for which only a portion of the offered quantity might be accepted.

Opening Rule: A new tender can be submitted only in the first iteration. After the first iteration, the only valid tenders are those submitted in the first iteration or validly revised subsequently that have not been withdrawn.

Exclusion Rule: An active step on a supply tender becomes frozen after the current iteration if its offered price is not validly revised to improve the previous clearing price, and in the previous iteration its offered price was above this clearing price – called its Activation Price. A frozen step cannot be revised. A frozen step becomes active again after an iteration in which the clearing price is higher than its Activation Price.

Revision Rule: An active step can be divided into two active steps with the same offered price. An active step can be revised only by offering a lower price that improves the previous clearing price. That is, the revised step must offer a new price for the same quantity interval that is less than the previously offered price, and also less than the previous clearing price by at least the specified price decrement.

Withdrawal Rule: After each iteration except the last, each supplier has the option to withdraw a tender entirely and irrevocably from any hourly market. The clearing prices are re-calculated after the withdrawal round. For the purposes of the Exclusion and Revision Rules and setting Activation Prices, these become the clearing prices for this iteration.

Closing Rule: All hourly markets close simultaneously. They close automatically after an iteration in which no tender is revised. Otherwise, before time expires, the final iteration is announced, and the results of the final iteration become binding transactions at the final clearing price. After the final iteration, an accepted tender cannot be withdrawn and the supplier remains financially liable for delivery.

Appendix B

The Experimental Protocol

The experimental program was conducted entirely at Caltech's Laboratory for Experimental Economics and Political Science (EEPS) by Charles Plott. The lab is equipped with twenty-four 200 MHz HP Vectra PCs, of which one is reserved for administration of the tests. The software was constructed by H.Y. Lee from proprietary modules in C++ previously developed and owned by the lab. The software mimicked the essential features of the PX Protocol and the "standard" activity rules in Appendix A. The subjects were Caltech students solicited via an email broadcast announcement; their remuneration consisted solely of their trading profits – or $20 if the software crashed. New subjects received only verbal briefings on the procedural rules; most had participated in other experiments at the lab and therefore they were generally familiar with how such experiments are conducted. There were no formal de-briefings after the experiments. Testing sessions were typically conducted two or three times per day for one or two hours over two weeks.

In a typical session each of four to twelve subjects was assigned the role of a supplier or a demander. If the role was supplier, say, then the subject received a two-column list of the supply cost for each of a discrete set of quantities sold from its portfolio. The multi-market experiments assumed that this supply function is the same in each market, except for a separate fixed start-up cost in those sessions with this feature, whereas demand was different in each of two or three (or four) markets corresponding to offpeak, peak, and shoulder periods. Some sessions used the inelastic summer demand schedule and the eight representative portfolios prepared by London Economics to simulate the California mix. Testing in the early phase used simple linear schedules, varying the slopes to establish the sensitivity of the outcomes to the demand and supply elasticities. The total-energy tests were conducted by allowing such a supplier to allocate its total supply among the three markets.

The display on a subject's PC screen showed a list of the markets, each with the previous iteration's clearing price and tentatively accepted quantity, and the time remaining until the close of the current iteration. By clicking on a market, the subject obtained a smaller screen that displayed rows in which revised offers could be entered. Each step of a current tender could be divided into two steps at the same price by clicking on a button. Each revision was automatically time stamped and then checked to ensure conformity with the Exclusion and Revision Rules, and to ensure that the revised schedule was non-decreasing. After the close of an iteration, the markets' clearing prices were computed and displayed to the subjects, along with the quantities tentatively accepted from that subject. In the case of multiple offers at the same clearing price the Rationing Rule accorded priority to offers with the earliest time stamps. The clearing price was computed as the lowest price for which supply exceeded demand, or if demand and supply could be equalized, then it was the greater of the highest accepted supply offer and the highest rejected demand bid.

The auction administrator's screen showed a complete summary of the status of the auction. After the close of the auction, summary results were stored in a tabular form.

[*] This and other reports from this project can be downloaded from http://www.energyonline.com/wepex/reports/reports2.html