Evaluating Demand Response Programs and Tariffs: Looking for Effective Synergies

Steven Braithwait, Christensen Associates, Madison, WI

ABSTRACT

Industry experts agree that increasing the price responsiveness of consumer demand for electricity is a crucial element of efficient power market design. Increased price responsiveness may be obtained through mechanisms such as *dynamic pricing* and *demand response* programs. However, some controversy has developed over the design of DR programs, including whether they necessarily imply public funding, and therefore careful evaluation. This paper examines the impact of program design on the economic benefits of DR programs, including benefits and costs to non-participants, and the evaluation needs implied. In particular, it illustrates a market-based DR program design that operates like dynamic pricing, and thus requires no non-participant funding.

Introduction

Interest in demand response (DR) programs appears to have arisen largely from concern about frequent wholesale price spikes. It has been recognized that price spikes are exacerbated by the fact that few consumers see retail prices that reflect wholesale cost conditions, thus making their short-run demand for electricity highly inelastic. If some consumers were shown the true cost of power, they might be willing to cut back on a portion of their less valuable consumption.

Load reductions during periods of high wholesale costs have the potential to cut wholesale price spikes substantially. They can also produce cost savings in excess of the value that consumers forego from curtailing load. These potential cost savings include both the avoided *operating costs* of the most costly generating units, and some fraction of the *cost of new generating capacity* that might otherwise be built to meet the higher level of unresponsive demand.

This problem of disconnection between wholesale and retail power markets has been recognized for some time. Power markets are typically characterized by a unique combination of non-storability of the supply of electricity, and unresponsive short-run demand due to a regulatory tradition of average-cost pricing. As a result, hourly *wholesale electricity costs* frequently diverge from fixed *retail prices*. Furthermore, the distribution of hourly costs is highly asymmetric, with many low-cost hours and relatively few very high-cost hours.

Methods for Increasing Price Responsiveness

Consumers can be exposed to wholesale market costs either through the retail prices that they pay, or by having the opportunity to effectively "sell" a portion of their otherwise planned consumption back into the wholesale market through a DR program. Perhaps the most straightforward approach to increasing consumers' price responsiveness is for utilities or load serving entities (LSEs) to offer some form of *dynamic pricing*, such as real-time pricing (RTP) or critical peak pricing (CPP) to at least some of their customers. Dynamic prices change on a daily or hourly basis to reflect changes in wholesale power costs, and provide consumers with a natural economic incentive to reduce consumption when power costs and prices are high. One of the most successful examples of dynamic pricing in recent years is Georgia Power Company's RTP program, in which more than 1,500 industrial and commercial customers comprising 5,000 MW of peak demand subscribe to face hourly prices that are announced on either a day-ahead or hour-ahead basis. However, utilities and state regulators in general have been slow to implement dynamic pricing to any great extent.

Partly as a result of this inaction at the state level, most attention has recently focused on *DR programs* offered through regional ISOs, such as PJM, the New York ISO, and ISO New England. In these programs, LSEs bid load reductions relative to a baseline load level, at particular prices. If the ISO accepts the bids, then they are scheduled into the relevant day-ahead or real-time market along with generators, and the LSE is paid the market price for the load reduction. The LSE makes analogous arrangements with its customers, who provide the necessary load curtailments in return for a payment of some portion of the wholesale market price (*e.g.*, paying customers 90% of the market price appears typical in New York). In addition, some DR programs include an *incentive payment* to the LSE for DR load reductions as an incentive to participate in the program (see discussion below).[1]

Evaluating DR Programs

Traditional program evaluation methods can be used to assess several aspects of a DR program. First, there is the key question of how much load reduction is provided by customers participating in the program during a particular period of operation (*e.g.*, the four to six hours on a day in which wholesale prices reached a given level). Measuring this amount of load reduction involves the typical program evaluation problem of estimating consumers' *baseline loads*, or the level of hourly consumption that they would typically use on such days in the absence of the DR program. In addition, there is the issue of evaluating the cost effectiveness of the program – what are the cost savings produced by the load reductions, and how do they compare to the costs incurred by consumers and their suppliers to accomplish them. Finally, there are issues of how the program was implemented and marketed, and how consumers reacted to the program.

Extensive evaluations have been conducted of the performance of the NYISO DR programs during the summers of 2001 and 2002. Estimates of load reductions, program benefits and costs, and factors that affect customer load response have been examined in some detail.[2] Considerable valuable information on the operation and effects of the programs has been produced. However, a review of the design of the day-ahead DR program in particular, and the estimated benefits, suggests that several improvements are possible. In fact, a key thesis of this paper is that a well-designed DR program can mimic the performance of a market-based dynamic pricing program, be self-financing, and thus require no ratepayer funding, or publicly-financed program evaluation. LSEs will need to conduct their own evaluations to develop methods for effectively marketing the programs, and for accurately forecasting their customers' price responsiveness and incorporating their load reductions efficiently into their load schedules.

Evaluation Methods

This paper demonstrates through a three-step process how market-based DR programs may be designed to operate similarly to a dynamic retail-pricing program. It then evaluates the benefits and costs of such programs, and finally examines implications of these designs for program evaluation. First, we examine the traditional case of fixed retail prices, and illustrate the potential for increasing economic benefits through more efficient retail pricing. Second, we demonstrate how one method of dynamic pricing – two-part RTP – provides consumers with economic incentives to respond to market-based prices, while at the same time protecting them from the financial risk of price volatility. Finally, we illustrate how a DR program may be designed to operate like a part-time dynamic pricing program,

---

[1] An historical prototype for these DR programs was the informal "buy-back" programs offered by a few utilities in the Midwest after extremely high wholesale prices erupted in the late 1990s. In these programs, utilities offered to split the savings from avoiding expensive power purchases if certain customers would agree to curtail load for a few hours on very high-cost days.

[2] See Neenan Associates (2003).

provide the same type of incentives for load response, and produce opportunities for both customers and their LSEs to benefit from the program.

Economic Benefits from Price Responsive Demand

We illustrate the potential beneficial effects of dynamic pricing and DR programs using Figure 1, which represents conditions in a representative hour in the day-ahead energy market. The figure shows a steeply-sloping *supply curve* in the range of high load levels relative to available capacity. It also contains three alternative consumer *demand curves* – two vertical lines representing unresponsive demands under "normal" and "hot" summer conditions when consumers face fixed retail prices, and the sloping demand curve labeled Demand (hot), which represents price responsive loads in the presence of dynamic pricing or a DR program. Key elements of the figure are as follows:

- On a hot summer day *without price responsive loads*, consumer demand increases from Q_{normal} to Q_{spike}, causing wholesale prices to rise to WP_{spike}.
- The incremental *cost* of producing the last unit of power to meet demand under the unresponsive scenario is given by the distance from the horizontal axis to point *B* on the supply curve (which represents incremental power costs at different levels of demand). The incremental *value* to consumers of that increment of demand is shown by the height to point *A* on the aggregate demand curve (which represents consumers' incremental value of electricity at different levels of consumption), at the fixed retail price *P*.
- If some consumers face dynamic prices that reflect wholesale costs, or can offer load curtailments into the wholesale market through a DR program, then aggregate demands are shown by the sloping demand curve, total quantity demanded falls to Q_{hot}, and the wholesale market clears at point *E* at WP_{hot}.

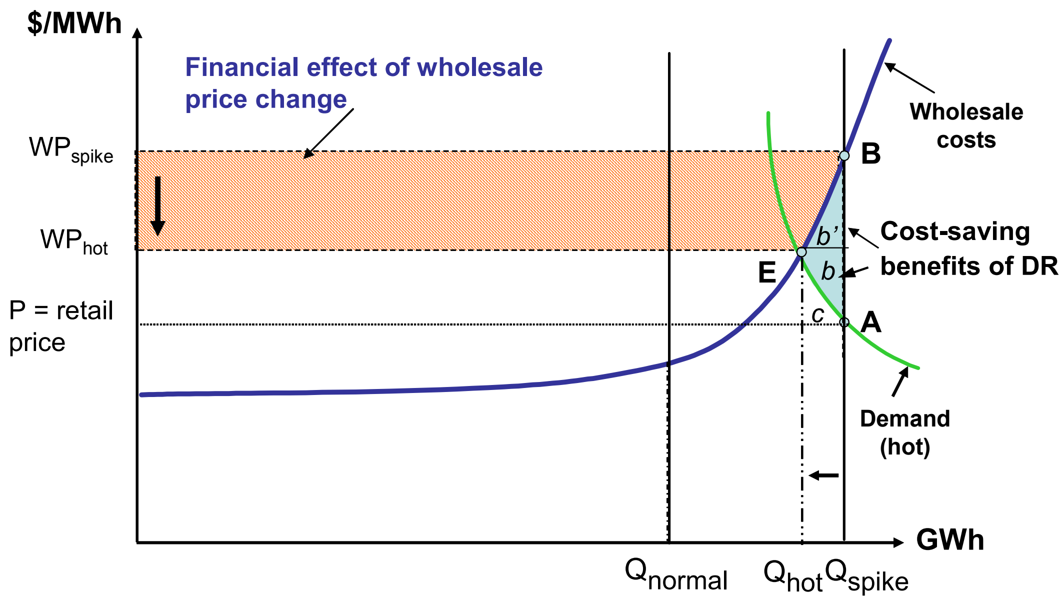

Figure 1. Changes in benefits and costs associated with price responsive load programs.

- Consumers' load curtailments allow the market to *avoid costs* equal to the area under the supply curve for the amount of the load curtailment. At the same time, consumers *forego value* from the electricity not consumed that is equal to the area under the demand curve for the amount of the curtailment. The indicated difference between those areas ($b + b'$) represents the net *cost-saving benefits* of price responsive loads. Depending on how dynamic pricing or DR programs are designed, those benefits can be *shared* between the consumers who curtail load, their energy suppliers who experience reduced supply costs, and potentially all consumers due to lower average power costs (*e.g.*, lower capital costs from avoiding the cost of generators needed to meet unresponsive demand).
- The *financial effects* of the wholesale price spike reduction, shown in the shaded horizontal rectangle, are often described as examples of the large potential benefits of DR programs. However, there are several problems with this interpretation. First, bill and revenue effects resulting from *price changes* are treated by economists as *transfer payments*, not changes in economic benefits, because they do not reflect changes in real economic resource cost or value, such as those described above that are caused by consumers' *load changes*.

  As noted above, wholesale price volatility is an inherent feature of electricity markets. High wholesale prices in particular serve an important role of signaling the value of generating capacity to encourage new investment when it is needed. In the short run, wholesale price volatility has little financial effect on most market participants because nearly all of them manage their price risk by either owning generation, or entering financial contracts to buy and sell energy at fixed prices. As a result, they are for the most part not directly injured financially by wholesale price spikes, nor helped by reductions in the price spikes. In the long term, wholesale prices reflect the capacity investment decisions made in response to historical prices. Thus, while DR programs may indeed serve to hold down short-term price spikes, *conditional on the existing generating capacity mix*, from levels they would otherwise reach, generators will take such reductions into account when making investment decisions for the future. As a result, short-term price effects cannot be assumed to hold into the future.

Dynamic Pricing – Two-Part RTP

Given the option, most consumers will choose not to face the volatility of wholesale power prices directly (*e.g.*, spot market pricing) without the financial protection provided by some type of price risk management mechanism. That is the primary reason that the only RTP programs that have proven successful to date have been *two-part* designs. These programs effectively offer customers a financial *contract for differences (CfD)*, which guarantees a fixed energy price for a fixed baseline level of usage. Under this design, customers pay market-based RTP prices for their entire load, but then also receive a financial adjustment to their bill, based on the CfD, that ensures that they pay no more than the guaranteed price for their baseline load. This design is also often described by the equivalent characterization that customers pay the fixed price for their baseline load, and then *pay* RTP prices for any load in excess of the baseline level, and receive *credits* at RTP prices for reductions below the baseline level.

We can illustrate how two-part RTP operates a period of high wholesale costs using Figure 1, and defining RTP customers' baseline load as Q_{spike}, and the contract price for the baseline load as *P*. In a period in which the market price rises to WP_{hot}, the RTP customers reduce consumption to Q_{hot}. Focusing on the portion of the two-part RTP bill related to the load reduction, customers first pay for that amount of load (as part of their baseline load) at *P*, and then receive a *credit* for the amount of the load reduction at the RTP price, WP_{hot}. The RTP consumers' net credit payment from the LSE thus equals the rectangular area ($b + c$). However, the consumers incur a cost due to the load curtailment equal to the foregone consumer surplus under the demand curve and above P, shown by c. This implies a net overall gain from RTP load response of the area b.

Note that while the LSE effectively pays customers for the load reduction, the LSE also avoids the cost of buying that amount in the wholesale market at WP_{hot}, and thus is no worse off. If the supply curve were horizontal in this range, then that would be the end of the story; the RTP consumers would receive all of the welfare gain from the load response.[3] However, given the rising supply curve in the figure, suppliers in aggregate achieve net cost savings equal to the area b'. Regulated utility LSEs would ultimately pass these cost savings on to all consumers in the form of lower average prices.

The above description of the design, operation, and net benefits effects of two-part RTP illustrates how dynamic pricing can create cost-saving benefits relative to the alternative of inefficient fixed retail tariffs, *without the need for financial incentive payments provided by non-participants*. It is only natural to ask – if two-part RTP can achieve this feat by giving customers dynamic pricing signals *all of the time*, can a DR program be designed to accomplish the same objective by allowing customers to bid load reductions at specified prices only *occasionally*, such as during the relatively infrequent periods of high wholesale costs? The answer is yes, as shown next.

Market-Based DR Program Design and Impacts

We now specify a market-based DR program, and analyze the resulting changes in benefits and costs. The program applies to consumers who face a fixed retail price, but have the opportunity to respond to and benefit from a price that reflects cost conditions in the wholesale market. The program has the following properties:[4]

- LSEs bid DR load reductions (relative to a baseline level of demand that they would otherwise schedule) into the day-ahead market at specific prices.
- If the LSE's bid price is less than the cost of the generators that would otherwise be scheduled to meet the total market load (including the LSE's baseline load), then its load reduction bid is accepted, and the LSE is paid the day-ahead wholesale market price for the scheduled load reduction.
- In the ISO financial settlement for the day-ahead market, the ISO *charges* the LSE the market price for the LSE's total scheduled *baseline* load (*i.e.*, including the amount of the DR load reduction), then *credits* the LSE with a DR payment at the market price for the amount of the scheduled load reduction. The net effect is that the LSE is responsible for purchasing only its *actual* scheduled load, net of the DR load reduction. [In addition, the ISO facilitates the clearing of any bilateral supply contracts that the LSE may have with generators at fixed prices.]
- To achieve the DR load reductions, LSEs must offer some incentive mechanism to entice their DR customers to curtail load. For example, the LSE and customer first agree on a method for calculating a baseline load that represents what the customer would otherwise have used. The customer pays its fixed retail price for the baseline level of usage, and the LSE credits the customer for any load curtailment below the baseline level that the customer promises to make, at a price tied to the wholesale market price (*e.g.* 80% of the market price).[5]

[3] This outcome suggests a variation on this design in which the supplier credits the RTP customer at something less than the market price for load reductions below the baseline level, as is the case in buy-back programs or a market-based DR program, and thus shares in some of the gain.

[4] This market-based DR program design is developed in greater detail, including the associated economic benefit and cost impacts in Braithwait (2003).

[5] Our understanding is that the New York distribution utilities do not charge their day-ahead DR customers for the amount of the load reduction at retail rates before crediting them for the load reduction at 90 percent of the market price.

Now review the changes in costs and benefits to each of the parties. For simplicity in using Figure 1 to illustrate the results, we assume that the LSE passes on the full market price to the customer for any load curtailments. A more realistic case would have the LSE pay some fraction of the market price, retaining a portion to cover its costs. The effects on each party are as follows:

- The *customer* receives the DR payment (equal to WP_{hot}) for his load curtailment, after paying the retail price *P* for that same amount. The customer's net DR payment equals the area ($b + c$). Since the customer voluntarily offers this curtailment, the net payment logically exceeds his forgone value, or cost of the curtailment (shown by *c,* the area under the demand curve and above *P*), resulting in a net DR benefit equal to the area *b*.
- The *LSE's* transaction with the ISO is a wash; its DR payment from the ISO completely offsets the settlement charge at the market price for the amount of the load reduction. For a given market price, he is able to avoid the cost of buying the amount of the load reduction at that price by paying the customer the market price to curtail. To the extent that the LSE actually splits the DR payment with his customers, he is able to retain a portion of the area *b* as his net gain.
- Suppliers in aggregate avoid energy costs in the amount of area *b'* as a result of the customers' load curtailments. Those LSEs that are regulated utilities presumably pass on these cost savings to all consumers. The combined *increase in net economic benefits* to consumers and suppliers amounts to ($b + b'$), the same amount as in the dynamic pricing example.
- Finally, the *ISO* is indifferent between paying a generator the market price for supplying power to meet the LSE's total baseline load, and paying the LSE the market price for a DR load reduction from that baseline level. In either case, the LSE pays the ISO the market price for the amount of the load reduction up to the baseline level. No uplift charges to non-participants are needed.

The key feature of the market-based DR program outlined above is that both DR payments – the one from the ISO to the LSE, and the payment from the LSE to the customer are fully covered by offsetting revenues or avoided costs, implying that the payments are *self-financing,* as with two-part RTP or other dynamic pricing methods. One possibly important difference between RTP and the market-based DR program is the method for estimating the consumer's baseline load. Under RTP, the baseline load is set in advance, thus alleviating the need for calculating moving-average baseline loads each day under most DR programs.

Comparison with Current DR Program Design

The design of some current ISO DR programs differs in an important way from the above market-based design. For example, the NYISO economic day-ahead DR program (DADRP) operates much like the above design, but with one crucial difference – in addition to the self-financing DR payment from the ISO to the LSE for the amount of the load reduction, the program also includes an *incentive payment* from the ISO to the LSE. The amount of this incentive, or rebate payment equals the market price times the amount of the load reduction. Thus, in effect the ISO pays the LSEs *two times the market price* for load reductions relative to their scheduled baseline load.[6] The LSE in turn passes on 90% of the market price to the DR customer.

Implications for Program Evaluation

Energy program evaluations are typically undertaken of programs that are sponsored by regulated utilities or state government agencies using funding obtained through surcharges in the rates

[6] It is interesting to note that the NYISO does not offer this additional incentive payment for customer load curtailments that are achieved through operation of a local distributed generator.

paid by most or all energy consumers. The evaluations are designed to ensure that the programs are undertaken in an efficient manner, using industry standard practices, and that the resulting changes in energy consumption are quantifiable and the programs are cost effective. To the extent that demand response programs involve the same type of public funding, they should be subject to comprehensive evaluations to measure the amount of load reductions that are produced, the changes in benefits and costs to various parties affected by the programs, and how effectively the programs have been implemented. However, as demonstrated in this paper, the very power market conditions that suggest the need for DR programs also create market-based incentives for self-financing DR programs that provide win-win opportunities for participants and LSEs. For programs of that type, no formal program evaluations should be needed.

Some types of LSE analysis will need to be undertaken for standard business reasons, including the following:

- Assessing the operating costs and the effectiveness of the marketing methods of the program,
- Developing and validating methods for estimating the baseline loads that are used for calculating the load reductions for which consumers are paid,
- Forecasting the load reductions that LSEs can expect at various price levels so that they can schedule loads accurately in day-ahead and real-time markets.

References

S. Braithwait, "Demand Response is Important—But Let's Not Oversell (Or Over-Price)," *The Electricity Journal*, June 2003.

Neenan and Associates and CERTS, "How and Why Customers Respond to Electricity Prices: A Study of NYISO and NYSERDA PRL Programs," January 2003.