NERA
ECONOMIC CONSULTING

Effective Use of Demand Side Resources: The Continued Need for Availability Payments

October 23, 2013

Project Team

Jonathan Falk, Vice President

Noah Kaufman, Consultant

Stephen Buryk, Analyst

NERA Economic Consulting
1166 Avenue of the Americas
New York, New York 10036
Tel: 1 (212) 345-3000 Fax: 1 (212) 345-4650
www.nera.com

Executive Summary

We have prepared this White Paper to clearly state the case for availability payments to demand side resources in electric markets. By *availability payments* we simply mean that there are promises of payments (direct financial or in-kind) to electricity customers in return for reductions of electric demand when requested. Availability payments to customers are a part of many electricity markets, so in some sense the proposition that availability payments make sense should not be controversial. What makes them controversial is the assertion by some economic theorists of electric markets that so-called *dynamic pricing*, in which customers pay prices for electricity which reflect the marginal economic cost of producing power, would render availability payments, or indeed, most demand-side programs, superfluous. These arguments are then recycled by advocates for generators who see demand-side competition as unwanted.

As this paper will demonstrate, the case for availability payments flows from the simple fact that dispatchable demand-side resources form a useful tool for electricity control room operators whose vital role is to keep supply and demand in balance. Their usefulness stems both from the fact that an additional source of control raises reliability and lowers the probability of blackouts, but also from the fact that these resources may well be more economic than incremental generation to provide load/supply balance.

While we could certainly imagine demand-side resources who provided supply/demand balancing services without availability payments, both theory and experience demonstrate that it would be unwise to try and implement such a scheme. The critical theoretical reason involves what has now become known as *behavioral economics*. Since these demand-side programs are all voluntary (and it would make no sense for them to be involuntary, since involuntary load shedding is what we are trying to avoid) we must find that set of customers who are willing to shed load. Finding these customers and convincing them that participation in these programs is in their interest is no easy task. Availability payments make the task of recruiting demand-side resources much easier, by overcoming well-accepted cognitive biases which make it difficult for customers of all types (residential, commercial and industrial) to uncover their best interests. We discuss these biases in some detail and show how, in theory, availability payments help overcome them.

The other theoretical underpinning for availability payments comes by analogy with capacity markets for competitive generators. Capacity markets, in which generators receive payments to make *their* capacity available to control room operators, were similarly deemed theoretically superfluous by economists who felt that simply paying generators the marginal value of their energy would suffice to yield sufficient capacity to keep the lights on. In the three large Eastern US markets (PJM, NYISO and ISO-NE) there is now general consensus that merely hoping that energy prices alone will incentivize market participation is not good enough. Even without a formal capacity market, many jurisdictions go through a planning and procurement phase which promises payments upfront in advance of any exposure to energy prices. Demand-side resources may well participate in these processes directly, in which case their availability payments are simply the same sort of payments generators receive, for the same reasons. Even in traditional vertically integrated electric systems, though, demand-side resource participation will, in general, be both (a) valuable to the system; and (b) best recruited through the use of availability payments.

The case for these payments is not merely theoretical, however. We have created a database of participation rates for over 900 programs in the United States. Programs which make availability payments have much higher participation rates than programs which depend on consumers to use pure moment-to-moment economic calculus to determine participation. Unsurprisingly, electricity customers are unwilling (and often unable) to put in the hard work of optimizing their electricity consumption from moment to moment; dynamic pricing schemes simply *assume* that consumers can and will do so.

Demand-side resources represent a potential source of cost-saving and reliability enhancement for *any* electric system. Successful recruitment of demand-side resources is substantially enhanced by the presence of availability payments.

Effective Use of Demand Side Resources: The Continued Need for Availability Payments

This White Paper explains why availability payments to dispatchable demand resources should be an important component of virtually any electricity market. The case for availability payments does not require centralized pooling markets, vertically integrated utilities, or installed resource procurement curves. The rationales continue to apply even in the presence of dynamic real-time pricing in which prices paid by loads continuously reflect short-run conditions on the network. This last conclusion may be surprising to some, but it should be no more surprising than the now commonplace notion that capacity markets for generators can form a useful adjunct to energy-only electricity markets.

Markets will exist whenever there are gains from trade that exceed the transaction costs of market formation[1]. In many contexts, such markets form spontaneously; in heavily regulated markets like electricity, however there is an extra set of forces with which policy makers must contend: stakeholders whose parochial economic interests might be harmed even as welfare is enhanced overall[2]. This problem is not unique to electricity markets; this problem is routinely faced in antitrust analysis, for example. As this document will detail, availability payment programs which are welfare-enhancing should be the default for electricity markets. The burden of proof should lie on any attempt to not include availability payments, rather than the current burden of proof which seems to go in the other direction, largely based on a constrained and overly academic construct of economic theory.

I. Problem Definition

Electricity systems everywhere operate on real-time controls. Within physical limits defined by Kirchhoff’s Laws, electricity systems must change the level of output of a number of electric generating units, change the configuration of the network dynamically, and/or change the level of electricity demanded by loads in real time so as to balance, within certain limits, electricity flowing on the network. Failure to do so can lead to various forms of interruption, including catastrophic cascading failure.

As a general proposition, the more points of control that the system operators have, the better their capability of keeping the system stable. This proposition is not ironclad, however, since it is always possible that unexpected interactions of controls can actually

---

1 Or, to state it more simply, if there are trades that make both sides better off and the cost of people finding each other is low enough, people will find each other to make those trades. There may be middlemen facilitating these trades, in which case their costs have to be low relative to the gains to the trading parties.

2 This includes even decentralized pooled electric markets, none of which functions anywhere in the world without an extensive regulatory superstructure

make the system more unstable – that fewer controls could be preferred to more. But the clear trend in the past fifty years has been to incorporate more and more sophisticated simulation and real-time automated controls, combined with an increase in the number of points to control.

Controlling loads is therefore a desirable goal. But in order to ensure enhanced welfare, loads must *volunteer* to be controlled. This non-coercion principle, while underlying all markets, is particularly relevant in electricity markets with a substantial degree of heterogeneity in consumer surplus[3]. Because all consumers are not alike, rules that rely on the behavior of an idealized consumer, or an average consumer, however well-measured, will almost surely leave some better off than others. In the absence of fairly precise knowledge about the *distribution* of preferences (which we do not have and have no reasonable prospect of procuring) we cannot confidently make any statement about welfare improvements. Because all consumers are not alike, *"one size fits all"* solutions are unlikely to work very well. That said, administrative simplicity probably requires a fairly small set of relevant programs. So the choice of programs to have should address not merely theoretical purity, but also empirical effectiveness.

One means to secure voluntary co-operation is an availability payment program where a load receives an upfront payment in return for some measure of control over their demand for electricity. There are good reasons why a program like this ought to be in the toolkit of every electric system, no matter how it is organized. This proposition is noncontroversial and such programs, in various forms, are part of most electric systems around the world. But some have argued that the potential rise in so-called dynamic pricing, under which loads directly pay some measure of marginal cost on a very short-term basis, makes such availability payment programs superfluous, or worse, inefficient. These arguments are incorrect, as will be discussed further below.

II. Paper Organization

First, we discuss the general organization of electric markets, which have many elements which are independent of the market structure. In particular, we focus on what we call the *control room problem*, which is whether the control room (which all electric systems have) has sufficient tools at its disposal to avert blackouts.

The next section describes one subset of the potential tools at the control room's disposal: demand-side programs. These programs have one thing in common: they seek to address the control room's problem by focusing on reducing demand rather than by increasing supply or reconfiguring the network.

[3] Consumer surplus is the difference between the value of a given service and what a consumer pays for it.

After this, we discuss what makes the demand-side choose to participate in a demand-side program. In classical economic thinking, each economic actor takes whatever actions are in his or her own best interests. This premise requires economic actors to understand what their best interests are. We explore in some detail the problems with assuming that economic actors possess this understanding. To those whose first language is not economics, this discussion may seem belabored. The proposition that consumers may not always take the best actions, or that the concept "best" may be an ambiguous one, may seem sufficiently obvious to many policymakers that the emphasis on this point may seem puzzling, but in fact the particular well-understood caveats to consumer rationality discussed here form the basis of our conclusion. Consumers are neither so hyper rational as to make availability payments unnecessary, nor are they so irrational as to fail to respond to directly understandable incentives. As the title of Duke University economist Dan Ariely's popular 2008 book explains it, they are "Predictably Irrational."

The next section explains why availability payments, either direct or indirect, serve to overcome the specific barriers to participation that are outlined in the previous section.

The last two sections discuss the empirical data which underlie and inform the assertions of the section on availability payments and barriers to uptake. The first of those sections relate to the 1970s and 1980s era debates over why consumers don't make enough energy-savings investment, i.e., why they do not make choices that appear to be in their best interests in economic programs. The second section looks at a contemporary database and shows that despite enormous changes in markets and the availability of so-called economic programs, that involve exposure to dynamic prices, participation in economic programs remains low compared to those for programs which make availability payments, either direct or in-kind.

III. Electric Markets

IESO Control Room, Ontario

Electric markets are organized so that loads can consume as much power as they choose to consume at any time with no notice to any central dispatching authority. The central dispatcher (which we will call a "control room" so as to make this discussion as broadly applicable as possible, encompassing ISO-run competitive markets, government-run markets, and traditional regulated utility markets) monitors and anticipates load and ensures that supply and demand are balanced at all times. This endeavor is extremely complex, because it must take account of the transmission system, the distribution system, transient events, sudden interruptions, voltage limits, thermal limits, losses, and a host of other concerns. The modern control room represents a formidable triumph of engineering prowess and human management.

Like all human endeavors, though, there are occasional failures. These failures are very costly relative to the cost of the service provided. To within an order of magnitude, electric service costs consumers around $100/MWh consumed[4]. Involuntary cutoffs of service, however, are orders of magnitude more expensive: a recent worldwide synthesis[5] suggested that the value of lost load varies from between $9,000-$45,000/MWh, hundreds of times the price paid, on average.

Thus, although power costs perhaps $50/MWh to produce,[6] a failure to produce power might cost $45,000/MWh or more. Given this universal reality, all electric systems run very conservatively, planning for substantially more capacity than is needed at any given time and making sure the system is robust to single transmission outages. But this conservatism is expensive, and any additional tools at the disposal of the control room to keep the lights on and which also reduce that cost will add to system efficiency.

Uncertainty in supply/demand balance has components on the demand side as well. Electric systems are structured so that consumers can demand as much or as little power as they like from moment to moment without informing the control room. For the most part, they are not even required to give indicative signals as to how much they will consume in a future period. This continuous option to consume power is an extremely important feature of electric systems; without it, the value of consumption would fall sharply. This then requires demand-side measures to work on a *voluntary* basis. Consumers who participate in demand-side programs are giving up flexibility. Since this flexibility is such a large component of value, the programs must find consumers who value that flexibility the least. This is the environment in which demand-side programs have to operate.

While it is clearly possible to compensate generators for extra capacity they bring on line to match either unanticipated demands or supply losses elsewhere in the system, it is more difficult to compensate loads for performing the same supply-demand balancing service simply because the amounts they would have consumed otherwise are more difficult to measure. Various approaches to get around this problem directly have been proposed (*e.g.* so-called “buy the baseline” programs in which demands commit to purchase an amount of electricity in advance) but those programs have not been adopted for a simple reason: they so inhibit the existing right of the electricity consumer to

[4] The most recent available US figure is $96.70. The US is in the middle of the pack internationally. See EIA, http://www.eia.gov/electricity/monthly/epm_table_grapher.cfm?t=epmt_5_6_a and NUS Consulting Group, 2011-2012 International Electricity & Natural Gas Report & Price Survey, June 2012

[5] http://www.ercot.com/content/news/presentations/2013/ERCOTValueofLostLoad-LiteratureReviewandMacroeconomicAnalysi.pdf

[6] The rest of the price is the fixed costs of running the grid.

consume as flexibly as he wishes that such programs are unlikely ever to gain consumer acceptance. This will be discussed in further detail below.

IV. Demand-Side Programs

Most electric systems have developed so-called demand-side programs, in which customers are given some sort of enticement to conserve on their demands for power. The genesis for such programs in the 1980s (long before there were any restructured wholesale electric markets) was a rise in the cost of building new generation. Rather than spreading that cost across all customers, the observation was made that incentivizing those customers whose peak consumption was not very valuable to reduce their peak consumption. If it was cheaper to reduce peak demand than to build the generation necessary to serve it, all consumers would be made better off: those participating in the program were better off because the voluntary nature of the program ensured that no one reduced consumption with some incentive which more than compensated them, and those not participating are better because their cost of electricity, even after paying the incentives, were lower than they would have been.

Other programs were developed to arbitrage between the traditionally fixed retail rates paid by customers and the very large marginal cost differences between peak and off peak power. For a number of reasons, consumers tended to be charged a fixed price per kilowatt-hour consumed even though power at peak times might cost twice the average price and power at off peak times might cost half the average. Inducing customers to pay for power in ways that more closely reflected marginal costs again had the potential of saving money for all consumers, both those who chose time-varying rates and those who did not, as average prices fell when those in the program shifter their demands from peak to off peak, or simply chose to consume less at peak times.

These programs have continued in deregulated markets: sometimes they are run by the central grid operator; often they are run by a load serving entity responsible for paying for the electricity taken from the grid by all its customers.

We will survey the most prominent methods below, but a critical point about all of them is that they are *voluntary*. This is not an accident. Only voluntary programs have the ability to make both those who select the program and those who don't better off. For those with an intense interest in consuming peak power (either at the system peak or in some broader peak period) programs which induce them to consume less at the peak will have to make larger payments; for many, those payments may be so large as to not yield any savings at all to the system. The programs work by establishing a level of inducement which, if taken, will guarantee direct cost savings to those not taking the plan. Then the fact that the plan is voluntary ensures that those taking it are better off as well.

Programs can be usefully categorized along two axes: dispatchable programs versus non dispatchable programs, and programs which provide capacity savings and those which provide energy savings. *Capacity* products enhance system reliability; they bring supply and demand into closer balance and thereby improve system reliability. *Energy* products reduce the total price paid for energy. Most products have both energy and capacity savings components, but generally one effect predominates. For example, a program which reduces load at the time of the system peak clearly has capacity benefits (since more capacity would have had to have been built to provide the same level of realized reliability) but there may be a minor amount of energy savings as well, even if only for a few hours of the year.

PROGRAMS WHICH REDUCE LOAD ON A DISPATCHABLE BASIS

Some programs allow the control room to issue commands to lower demand by a predictable number of MW; these programs represent dispatchable demand resources. The critical issue here is that the change in demand is determined by the control room, not the consumer. By giving the control room the power to request demand reductions, their tool set for running the system is increased.

- Interruptible Rates: consumers who choose interruptible rates pay lower prices for electricity at all times relative to those who do not, but agree to have their power cut off at times of system emergency. Plans vary in the number of possible interruptions, the duration of those interruptions, and the amount of warning provided. Some also make payments for interruptions. Interruptible rates form a classic capacity product. The interruptions are generally made only to enhance reliability, not to reduce price paid by all consumers during the interruption.

- Direct Load Control: Direct Load Control (DLC) is a partial version of interruptible rates. The control room has the ability to directly control a subset of some load’s consumption. An example would be air conditioning loads. Under DLC, the control room has the capability to cycle air conditioning on and off and thereby reduce total air conditioning load over all customers. Some programs enable the consumer to override the control under limited circumstances. DLC is again usually (but not exclusively) a capacity product, invoked when system reliability is threatened[7], thereby displacing generation which would otherwise had to have been built to provide equivalent reliability.

- Curtailment Service: Under curtailment service, a set of aggregated loads are compensated to reduce load below expected levels at a signal from their control room. They then receive a payment. Curtailment differs from interruptibility and DLC in that curtailment is generally only partial – consumers can choose which demand to reduce to fulfill the request from the control room. Curtailment is also efficiently provided on an aggregate basis, in which a portfolio of customers agrees to make a portfolio reduction in

[7] Some DLC programs have an energy component as well, calling load control when energy prices are high. There is often overlap, of course, between high prices and emergency conditions.

demand in response to a signal. Individual customers might not curtail or reduce load at all – only the aggregate reduction matters.

- Capacity Resource: In return for an upfront payment, loads agree to make capacity available under specified conditions; they do so by reducing demand below a measured baseline. The difference between the baseline and actual load is a contribution to capacity. There may be a payment for each curtailment as well, making these programs overlap somewhat with the curtailment services listed above. The only real difference is that these resources are provided in an organized capacity market which may include generation resources as well.

  The upfront payment obliges participating customers to curtail when system reliability is threatened, providing an important, reliable tool for any control room operator, providing useful controls with predictable effects. The *control room problem* is to ensure the control room has sufficient tools in its disposal to avert blackouts – only these methods allow the control room to equate supply and demand.

  The *value* of these programs is twofold: most importantly, the presence of an additional control point reduces the probability of load-shedding events. Less importantly for the control room problem, but nonetheless of some value, is the fact that the resources obtained may equate supply and demand less expensively than by scheduling additional generation at short notice. This second source of value is sometimes treated as the only source of value for these programs: it is not.

  Participants in dispatchable demand-side programs give up two things: flexibility and, when called, some consumption. In return they are paid. The payment must include both a payment for the lost consumption and a payment for the lost flexibility if it is to compensate them adequately; any voluntary program must compensate adequately if it is to be sustainable. Much of the confusion over the efficiency and payment for demand resources has missed this fundamental point. The assumption of many discussions of the amount to pay for dispatchable savings has focused on the lost energy value alone. Worse still, it has focused not on the social value of that lost energy but on the cost borne by the consumer. What loads need to be paid for dispatchable services and what their marginal incentives to consume at the margin are fundamentally different things, though with substantial overlap in purpose. If some subset of customers is willing to sacrifice flexibility and energy at a lower price than others to accommodate system balancing, then paying that price saves social resources.

PROGRAMS THAT REDUCE LOAD ON A NON DISPATCHABLE BASIS

- Time of Use Pricing: In Time of Use Pricing, prices are time-differentiated, but fixed. Thus, prices will generally be higher during peak periods and lower at off-peak periods thereby over time encouraging customers to change the timing of their energy consumption. These programs may have minor capacity benefits, but they are largely focused on reducing the price of energy.

- Critical Peak Pricing: In critical peak pricing, prices are occasionally increased at times specified by the control room. In a typical CPP program, prices might increase by three- or four-times the normal peak price for a few hours during the afternoon of an extended heat wave. In return, the subscriber to a CPP program gets lower prices the rest of the year than they would in the standard pricing plan. CPP programs require effective notification to subscribers to be effective. However, note that customers are not *required* to reduce load, they can simply pay more for their energy. Further, participating consumers would have had to be consuming normally at the critical moments if the program didn't exist; otherwise, a ski resort that uses no summer power could simply use a CPP program to reduce its electric bill without actually contributing anything. The programs generate mostly capacity benefits.

- Net Metering: For those customers with on-site electricity sources that exceed on-site load, net metering allows excess power generation to be sold back to some entity which then can either use it to reduce net retail load or sell it back in the wholesale market. In some sense these really aren't demand-side programs at all, but simply represent generators whose net supply to the system is erratic and uncontrolled by the control room. They generally represent energy savings to the system, not capacity.

- Dynamic Pricing, also called Real-Time Pricing: Dynamic Pricing means charging loads the marginal cost of power[8]. Because the marginal cost varies continuously, is not really achievable. Instead, pricing varies over some smaller increment of time – as little as five minutes or as much as an hour. Usage is metered at fine enough intervals to calculate quantity taken during the period when each price applies. So long as the energy price paid reflects both the cost of energy and the value of capacity, dynamic pricing can be expected to have both energy and capacity benefits. As will be discussed in more detail below, however, the energy price paid must reflect both energy and capacity benefits in order to have both benefits.

All of these non dispatchable programs are useful to control rooms. They all use price to curtail demand when prices are high and raise demand when prices are low. Because supply works in the opposite direction, i.e., it increases when prices are high and

---

8 *The marginal cost of power* is an idealized term. In fact, some calculated price intended to represent the marginal cost of power is used. This calculated marginal cost may be quite far from the "correct" marginal cost of power as thought of by economists. For example, it might exclude environmental costs which are unpriced. Or it might improperly reflect market power. Whether or not electricity supplied by demand responders better aligns with the "correct" marginal costs is a complicated question which will be ignored here. We will assume for this paper that the system operator is able to calculate some number which we will refer to as marginal cost (sometimes called LMP, or local marginal price) which is close enough to be regarded as a true social marginal cost.

decreases when prices are low, all of these programs are valuable in helping to balance supply and demand.

It is important to note, however, that they work through an intermediating mechanism: price. They are therefore, in general, *less* valuable to the control room to assist in the control of power systems than dispatchable programs, simply because they are less predictable. In addition, system emergencies can arise when power prices are low. For example, the sudden emergency tripping off line of a power plant during the fall season might create a temporary emergency that is outside of time-use or critical peak pricing periods.

A Digression on Dynamic Pricing

The idea of what we call dynamic pricing is not new, nor is the argument for its optimality. The first fundamental theorem of welfare economics[9] is that setting prices at marginal cost produces allocative efficiency: no individual values any good he consumes at more than the cost to produce it and production comes from the least-cost set of producers. The application of this theorem to electricity is found in Schweppe *et al.*, in their 1988 book, *Spot Pricing of Electricity.*

For this paradigm to actually function, however, loads must bid a demand curve just as suppliers bid a supply curve. Even without explicit bids, the system might function if demand curves are implicitly loaded into control systems which take price as an input from the central dispatcher and adjust loads to align marginal benefit with marginal cost. But neither universal load bidding nor universal appliance control are either politically or technologically feasible today. Technological feasibility might be near at hand, but political feasibility is not.

Numerous technical issues complicate the application of the first welfare theorem to electricity. First, electricity prices are usually calculated *ex post*, often days after the electricity is consumed. While a contemporaneous indicative price might be posted, such a price will greatly complicate the optimal consumption problem, which now is similar to a futures problem in which one agrees to take power and pay a price which is linked to the indicative price. Second, although blackouts are technically just another example of dynamic pricing (price → infinity and consumption → 0) the actual avoidance of blackouts is not a part of dynamic pricing programs. While the control room can *hope*, or even *expect*, that higher prices will induce enough load reduction to forestall involuntary load shedding, prices can easily move faster than they can be fed back to loads – Kirchhoff's laws operate at the speed of light; human controls work at a much slower rate.

[9] Statements of the First Fundamental welfare theorem can be found in any introductory economics textbook, e.g., Cowen and Tabarrok, *Modern Principles: Microeconomics*, Worth Publishers, 2011

In practice, setting the price of electricity to consumers to equal its marginal cost has another very difficult problem which has hindered its application. The marginal cost of electricity has two components: the marginal cost of the energy and the marginal reliability loss from the additional unit of consumption. The first component is relatively simple to estimate and is the natural price which will emerge from competitive bidding for power by generators and loads. The second component is very difficult to measure and generally stays unmeasured.

When the demand for power comes close to exceeding the supply of power, the probability of involuntary load shedding rises. As we saw above, involuntary load shedding is very costly to society: $9,000/MWh or more. But this reliability cost of increased consumption *cannot* be accurately measured through the interaction of supply and demand bids because of market power: when potential supply is close to demand, there is no competitive force which moderates the bids generators can make.

In practice, this problem is solved in one of two ways: the first is to add *scarcity pricing* to the price of power, an administratively determined adder to the energy price which reflects some administratively determined trigger of supply and demand threatening to get out of balance. The second method is to use a *capacity market* to administratively ensure (through commitment payments to loads and generators) that scarcity conditions are unlikely to occur.

The problems with both of these methods lies in the fact they are administrative and do not necessarily reflect the true underlying preferences of consumers. Thus, while the signals from scarcity prices and capacity markets can be incorporated into dynamic pricing schemes, the case for their economic efficiency can no longer be made so confidently.

Even ignoring this problem, however, the supposed panacea represented by dynamic pricing is most undermined by the simple issue of acceptance. Can dynamic pricing be made attractive enough for a sufficiently large number of customers to select it so that it actually works?

No matter how the electricity system is structured, *someone* bears the marginal cost of power. In the main US deregulated systems, load serving entities are charged LMP for power and those aggregate costs are then charged back to loads in a way intended to keep the load serving entity whole. In vertically integrated utilities, the aggregate costs of running the system include the marginal costs. The process of tariffing creates the average cost which consumers pay. Typically, then, loads are insulated from the consequences of their consumption decisions, paying, in effect, the average cost of power rather than the marginal cost.

Every consumer of electricity who is not exposed to the marginal cost of electricity has at least the potential to create *some* inefficiency. That fact does not imply, of course, that dynamic rates should be mandatory[10]. But it does mean that without some sufficiently large core of elastic consumers on dynamic rates, the efficiency effects will perforce be small.

And when we look at *voluntary* acceptance rates for dynamic pricing programs, even their strongest proponents admit that they have thus far been smaller than desired. Even as advanced metering has gone from under 1 percent of customers in 2006 to over 20 percent of customers in 2012[11], and is forecast to double again by 2015, a recent survey of industry experts shows even the most optimistic estimates have less than 10 percent penetration in the next ten years. Thus, even if dynamic pricing were a complete solution to the efficiency of electric systems, it is unrealistic to expect it to be implemented in bulk in any reasonable time period. So the more traditional style of demand response programs will be required for the foreseeable future.

This survey of demand response programs should make clear two things. *First*, demand-side programs encompass, optimally, a lot more than dynamic demand programs. Dynamic demand programs are not dispatchable; they require real-time metering infrastructure which many of the simpler programs do not require; they require sophisticated equipment on the customer's side if the signals are to be usefully reacted to. The other programs listed here have various characteristics which avoid these issues. While the dispatchable programs all require some kind of advanced metering in order to monitor compliance, their dispatchability assists in the control room problem in a way that dynamic demand programs cannot. Many of the other non-dispatchable programs may achieve large fractions of the benefits without the substantial cost of infrastructure improvements at the customer's premises. We should not expect any one program to dominate the others in terms of usefulness, either to the control room or to customers.

Second, all relevant demand-side programs are voluntary, and thus require some affirmative reason for demand-side entities to respond. Electricity consumption is voluntary everywhere. While changing the prices at which that electricity is sold may well have substantial efficiency benefits, it cannot address the dispatchability issue, which requires the consumer to give up consumption flexibility in advance. If these benefits are to be realized in a voluntary program, compensation must be paid to consumers to induce them to cede that flexibility.

---

[10] First, the costs of metering may exceed the savings for some customers. Second, some customers may have a sufficiently strong preference for fixed rates that they are willing to pay a premium above their efficiency losses for price certainty.

[11] FERC, *Assessment of Demand Response and Advanced Metering*, December 2012, Figure 2-1

A Digression on Customer Classes

Traditionally, electricity customers are divided into rate classes by type of use: at a minimum, there are rate classes for residential customers, small commercial customers and industrial customers. There are often more sub-classifications of these customers based on size or usage patterns. Many demand programs are offered to one class or another. In the empirical section, we will discuss particular customer classes and their rates of subscription for various services in more detail, but from a theoretical standpoint, we will make few distinctions between an individual household and a large industrial customer. That by no means implies that the process of signing one or the other up for a particular program is remotely similar. Arguments which are important to households may be irrelevant to commercial establishments, and *vice versa.* But from the standpoint of economic theory, residential customers and commercial and industrial customers must make choices from the options available to them and need to have well-defined criteria for making those choices. While we will occasionally contrast residential customers and commercial customers in the section below, we are uncertain of whether residential customers are more or less amenable to demand response programs, if properly designed. For concreteness we may describe choices as the sort of choices a residential or commercial customer might make, but this is not intended to limit the discussion in any way to one class or another.

V. What makes people sign up?

Now that it is established that we will continue to need traditional demand response programs, we can now turn to the central question: what will induce individual entities to participate in these markets?

We should point out that *any* demand-resource program has three criteria that it *must* satisfy:

(a) Participants in the program must be net beneficiaries of the service. It is highly unlikely that any program could survive which requires more sacrifice from the participant than benefit.

(b) However payments are structured, those paying for the service should be made better off as well. The reason for this should be equally clear – a program which does not benefit those who pay for it will not likely survive for long.

(c) The administrative costs of running the program should be far less than benefits obtained.

Any program that fulfills these three criteria is welfare-improving. Given that demand response programs have consistently been deemed cost-effective and have been approved by a host of Public Service Commissions, this Study is primarily concerned with the first criterion and largely assumes that programs which induce enough people to join can be configured to fulfill the second and third criteria as well.

A. Classical and Behavioral Economics

In this section we explore the economic theory of consumer choice and then apply it to the problem of demand-side participation in electric markets. While it is not the only possible theory of consumer choice, its persuasiveness rests on the notion that individuals tend to make choices in their best interests without being perfectly rational calculating machines. It should be noted that the basic economic theory taught for the last several centuries did explore the implications of humans as perfectly rational calculating machines and that assumption has proven remarkably successfully in answering a number of important questions, including ones in the economics of electricity.

We explore that theory here (and the recent challenges to it) because not to do so would cede too much to the advocates of dynamic pricing. In an ideal world of perfectly rational calculating machines, the scope for the sorts of traditional demand-side programs discussed above would be quite small. It is in the realm of classical economic theory that a case could be made for mandatory dynamic pricing of electricity with possible side-payments from the winners to the losers to make everyone better off. But in the real world such programs are impossible to implement. This means that in the real world customers have to be induced to voluntarily subscribe to programs which have reliability and/or energy benefits. And it is beyond dispute that participation or nonparticipation in these programs is not entirely determined by whether or not an individual household or business will or will not benefit by participation. It is here that the case for availability payments to induce participation in demand-side programs is made: availability payments take those customers most likely to be made better off by participation and actually get them to participate. How that happens will be the subject of the section following this one.

In classical economic theory, people[12] are assumed to make that set of choices which make them as well off as possible. From this insight, we get the *axiom of revealed preference*: if an individual had a choice between A, B and C, and chose A, then A must make him better off than B or C.

Even in classical economics, however, the definition of "well off" has a number of complications. First, of course, is the issue of knowledge. If an individual chose A over B and C, but had no idea that C even existed, the axiom of revealed preference will not

[12] To reiterate the point made in the digression on customer classes, we use the term "people" or "customers" in this section to refer interchangeably to residential customers and commercial and industrial customers. Where decision-making processes might be thought to be very different, caveats will be noted in the text. In point of fact, all decisions come down at the end to decisions by individuals, whether on behalf of an individual household or a multinational corporation.

apply. Second, there is calculational capability. If the choice between A, B and C requires a complicated process to determine which is best, then it is possible that the consumer will simply rely on rules of thumb, which may not always be accurate. Third, there is the issue of benefits over time. Many decisions have costs and benefits occurring at different points in time. The consumer must have some consistent method of assessing the net impact at the time the decision is to be made of the discounting of future costs and benefits. Fourth, there are preferences: each person will have a different value for electricity which varies from time to time based on their needs for power. Even holding the need for electricity constant, some consumers might value "green" power differently than other power. Fifth, there is the issue of uncertainty. The costs and benefits may be imprecisely known, or for that matter, imprecisely knowable. Consumers need to have some consistent measure of assessing their reactions to uncertain future outcomes before we can assess the reliability of their judgments.

Classical economics uses assumptions about consumer behavior to answer these five characteristics: the knowledge dimension, the calculational dimension, the preference dimension, the time dimension, and the uncertainty dimension. Once a consumer is specified using all of these characteristics, we can begin to make predictions about how he will behave.

In recent years, however, many researchers, so-called behaviorists, have begun to question the classical assumptions and suggest that consumers might be fundamentally inconsistent in each of these dimensions. Their actual behavior might vary from their "true" underlying knowledge, calculation ability, and preferences, and they might behave inconsistently when examining choices with different patterns of costs and benefits over time and different distributions of results under uncertainty. The set of results which has led researchers to these conclusions is too broad to detail here and a critical assessment of the accuracy of these observations fortunately is not necessary for our purposes. In this section, however, we will discuss each of these areas from both a classical and behaviorist perspective. The net results will convince the reader that there is no need to resolve these questions now, as both point to important reasons why programs that make up-front payments to induce participation are likely to be more effective than those which do not.

1. The Knowledge Dimension and the Search for Alternatives

In most contexts, classical economics assumes that people know what their choices are. However, if this assumption were true, we would not undertake so much effort to inform people. In some contexts (for example, job searches) we assume that people *do not* know all of their alternatives but systematically follow some plan to learn about a subset of their choices before choosing. In the context here, neither of these assumptions is likely. Consumers are certainly not uniformly knowledgeable about the various demand-side programs offered and few aggressively seek out that information.

That said, entities have made great strides in attempting to inform customers of demand response programs through advertising, bill inserts and other information sources. Although there undoubtedly are some consumers whose nonparticipation is simply based on unawareness of the existence of these programs, such lack of awareness is probably not the major impediment to widespread participation. More likely, failure to participate comes from other well-understood barriers such as doubt that the programs themselves can save the consumer much money, concern that the programs may prove burdensome, or outright confusion. Further, it is unrealistic to expect consumers to actively seek out such programs: electricity is for most households and businesses a fairly small part of total expenditures. Dynamic pricing, for its benefits, *assumes* that consumers understand the value of electricity from moment to moment and self-regulate their consumption, turning off light bulbs whose value falls short of the price paid, raising temperatures in the summer when the marginal cost of electricity exceeds the marginal comfort of air conditioning and in general never consuming a kWh of electricity whose costs exceeds its benefits. Whatever the problems that utilities and aggregators might have signing up clients for traditional demand services, those problems are dwarfed by the problems of forcing a consumer to understand his own consumption needs well enough to regulate his own consumption in reaction to price changes. This comprehension gap further undermines the case for dynamic pricing solutions.

Bounded Rationality

Choosing the best course of action is a mentally taxing activity. This is particularly daunting in circumstances where the choices are unfamiliar. Optimizing one's electric bill, especially by subjecting oneself to worse reliability, is simply a problem that few consumers have ever encountered. Although there certainly might be scope to *describe* the individual programs better, it is important to remember that the criterion here is not to convince people to take alternative tariffs, but rather it is to provide consumers with information that will allow them to choose whatever tariff best suits their needs, whether it happens to save energy or not.

In classical economics, consumers cannot really be fooled, or at least they cannot be fooled consistently. Behavioral economics makes the claim that consumers are in fact able to be consistently led to choices that do not optimize their true underlying preferences. Nevertheless, this assumption is beside the point: the important behavioral issue is that certain characteristics of programs may make consumers more likely to pick a particular program by economizing on the amount of thought that needs to go into the decision.

Learning how to behave in the world is a complicated endeavor. It is much easier to do whatever you did yesterday than to do something new. This fact underlies the ordinary marketing strategy that younger consumers are easier to market to than older consumers. This commonplace observation puts a burden on getting consumers to choose alternative tariffs, because one must not only present them with a new tariff, but make a compelling case to even consider it as part of a choice set. It clearly is not enough to simply make a program available, or even to make it available and prepare marketing materials to explain why the program might be worthwhile for some people; one must overcome natural consumer inertia as well.

2. Preferences

Electricity is a commodity with extraordinarily large consumer surplus. On average, electricity consumption is worth orders of magnitude more than is paid for it. As a result, people do not usually consider the value of individual uses of electricity – there is no need to do so very often. Indeed, while a large electric bill one month might induce a consumer to tighten his belt the next month, it is quite doubtful that one customer in a hundred knows that turning off a 100 Watt light bulb for ten hours saves exactly one kilowatt-hour, or has more than a vague notion of what a kilowatt-hour costs. And a light bulb, which actually states its consumption, is really the easiest case; the power draw of electronics or appliances is mysterious to most consumers. This is unsurprising, because they have no real way to disaggregate their consumption into the component pieces.

Even where a consumer has some notion of the actual cost of an appliance, he has little concept of methods for varying the consumption of that appliance and maintaining as much utility as possible. Few consumers would be willing to unplug their refrigerators on the hottest days or forego air conditioning no matter how much money it seemed to save. Rapidly cycling these appliances to maintain most of their utility while saving one-third of their consumption requires a consistent notion of the value of the service provided to be salient to a consumer. And until they have tried it, most consumers have no idea whether the resulting price-quality bundle is a good trade-off or not.

In classical economics, all consumers have an underlying preference for every use of electricity at all times. They can then decide whether or not, at the margin, a particular use of electricity has a value which exceeds its cost or not. Even to state this proposition is to highlight its unrealistic nature. People (and firms) simply do not have the detailed preferences required to make these decisions routinely. Combined with the consumer inertia described in the previous section, it is therefore unsurprising that they are unwilling to participate in these programs simply through the provision of information about the program itself.

Against this view, particularly in regards to dynamic pricing, it is argued that consumers need merely express their preferences in a control program which will take the relevant knowledge of current electric prices and the value of various uses and implement optimal choice. The usefulness of such programs generally is untested. It is highly unlikely that the "value of running the light bulb in the living room" can be contextualized in such a way that it can be usefully given to a program –the value varies with context and is probably not well understood even by the person making the decision.

Commercial and industrial customers, while perhaps better equipped to assess the bottom-line effects of particular uses of electricity, are no more capable than households in imagining scenarios of electricity use outside of their normal experience. Thus, while the examples above relate to residential customers, the same issues persist for all but the largest industrial customers. Even a large commercial establishment may have no one whose job it is to seek out the lowest cost electricity usage consistent with the company's mission.

3. Risk Preference

Even if consumers behaved perfectly according to well-founded optimization criteria, there is an independent reason why one might expect consumers to choose not to participate in a program that yielded them expected savings: risk preferences. Even if a program yielded gains on average, there might be many circumstances in which the particulars of the program created losses. To take an obvious example, most interruptible load programs leave little control to the customer as to when his load is curtailed. While the customer might *expect* gains, it is obviously possible for the customer to *realize* losses in any particular time period for which the lower rates do not fully compensate.

Further issues arise when we add the behavioral economics insight that most consumer attitude towards risk is inconsistent and incoherent. Some risks which are much lower than others are nonetheless treated as if they are much higher. Risks that are constant are not treated as constant in varied contexts[13].

Where expected value gains are eliminated because of risk avoidance, there is scope for programs that involve *risk transfer*. Risk transfer is the process whereby a party better able to assume risk, for a fee, assumes the risk that another party is unwilling to assume, enabling a welfare-improving transaction.

Classical economics assumes that consumers understand the payoffs of their choices in varied states of the world, but this sort of understanding almost surely exceeds the knowledge of individuals choosing electricity programs; indeed, it often exceeds anyone's actual knowledge. Where choices must be made in the absence of quantifiable certainty, people may become even more reticent to invest.

Dynamic pricing schemes, or for that matter any schemes which require extensive upfront investments, will be very difficult to sell when the upfront costs are highly salient and the benefits are subject to substantial uncertainty. Any set of characteristics which can lock in benefits will clearly lead to higher participation levels of customer uptake.

[13] See in particular Kahneman, Slovic and Tversky: *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, 1982.

4. Time Preference

Time preference represents another important feature of many choices. Benefits that accrue in the future carry less weight than benefits that are current or imminent. Each individual has his or her own rate of optimal time preference, and many households have rates of time preference which would make otherwise welfare improving choices uneconomic.

The issue of time preference in the investment in energy-saving technologies has a long history in economics[14]. Many investments that would appear to save money for individual households or businesses are nevertheless not made. While some have explained this outcome as caused by consumer inertia, ignorance, or capital constraints, the fact is that investments with long payback periods may simply not be optimal given the time preference of a household. This fact has led to a standard sort of arbitrage: another entity (one with a lower discount rate) makes the investment and reaps the majority of the rewards, sufficient to pay back the investment at the lower discount rate, and the residual reward belongs to the customer on whose behalf the investment was made. Many electric efficiency investments take this form, including utility programs that purchase (or greatly subsidize) energy-saving equipment (e.g., insulation), and the new breed of solar investments, such as solar panel leasing mechanisms[15].

Time preference issues, just as we saw with risk preference, will chronically deter participation. All programs require upfront expenditure of thought or dollars, usually both. The benefits, if any, will be spread out over the future.

VI. Why Availability Payments?

Now that we have looked at reasons why seemingly advantageous programs are not taken, we are in a position to see just why availability payments work so well. By availability payments, we mean some form of cash inducement to participate in a program, either one-time or periodic. Critically, however, the program participant realizes some or all of the benefits (in expectation) before bearing the costs.

The specific mechanism by which benefits are promised can vary. Some capacity payments might be made in advance of any consumption at all. It can be paid in arrears, so long as it is certain. It can be combined with specific payments for each time the resource is called on to provide resources. The mechanisms which work best may vary from customer to customer. But all of them provide a direct definite reward for ceding the

---

[14] The literature began with Jerry Hausman, “Individual Discount Rates and the purchase and utilization of energy-using durables.” Bell Journal of Economics, Volume 10, no 1. 1979, pp. 33-54. .This paper has been cited more than 900 times in the economics literature generally: http://scholar.google.com/scholar?cites=15686266506165200010&as_sdt=5,33&sciodt=0,33&hl=en

[15] http://www.solarcity.com/about-us.aspx

flexibility in consumption to which consumers are entitled. Thus, in any availability payment scheme, the customer is promised some form of benefit in advance of participation, even though the actual payment might be made later. The actual timing of the payments is not as material as the front-loading of the *promise* of payments.

The basic structure of capacity payments works to induce participation by promising certain benefits upfront. While this does not ensure that the benefits exceed the costs (customers can still make mistakes), a locking-in of the upside can induce plan participation. The next section of the paper takes the behavioral economics issues discussed in this section and shows how capacity payments can help spur participation.

If this notion is not sufficiently compelling, consider the alternative without availability payments. Many of these programs require substantial outlay of costs before any benefits can be realized. Some of these costs are technological (advanced metering, price signaling from the control room, automated on-premises response equipment) and some are, as detailed above, behavioral (inertia, cognitive issues). Asking a customer to bear *any* substantial costs before seeing *any* of the benefits clearly reduces the probability of participation. This is especially true when the customer is being asked to reliably *commit* for a period of time.

A. The Knowledge Space and the Search for Alternatives

When net benefits are frontloaded, the set of choices presented to the consumer is much easier. The customer is presented with a choice in which there are immediate net benefits, and the customer does no worse in subsequent periods than they would have done otherwise. There still remains the substantial burden of convincing the consumer of the benefits, but the other impediments to unfamiliar choice are greatly reduced.

B. Bounded Rationality

The complicated projection of costs and benefits and reduction to expected net present value is a daunting exercise for those versed in these calculations; for those whose only interaction with the electric system consists in installing and running equipment and paying bills, it is unsurprising that they are unwilling to even make an effort, and that even those motivated to make the effort are ill-equipped to make the decisions correctly. Upfront payments can simplify the decision greatly; risk adjustments are muted, time preference issues are simplified, and general uncertainty is reduced. Rather than calculate payback periods and discount rates, if a load can be guaranteed, some savings upfront followed by a path of behavior which would provide the conduit for those benefits, the rational decision making process is simplified.

C. Preferences

The process of actually deciding whether or not the costs of participation in a demand response program (potential cutoffs, cycling of equipment, programming of individual end-uses to respond to price, investment in energy-saving technologies, investments in backup power) is made easier by guaranteeing a level of return upfront. By itself, this process cannot guarantee that the consumer will be able to forecast his reaction to whatever inconveniences are put upon him by program participation, but without some inducement, it is quite unlikely that he will even be willing to contemplate the questions.

D. Risk Preference

Moving a substantial part of the net benefits up front clearly transfers at least some risk from the program participant to whichever entity is paying those benefits. The program participant *starts out ahead.* While this does not guarantee that the program participant will stay ahead, it clearly reduces risk.

E. Time Preference

Any program requiring the investment of capital, time or effort on the part of the program participant will require it upfront. If program participants have higher discount rates than society as a whole then putting the costs upfront with the benefits in the future will naturally cause too little social investment. Almost all programs that require capital investment subsidize the initial investment to get around this commonplace problem. Where the necessary upfront investment is consumer effort or the overcoming of consumer inertia, however, an upfront loading of the actual benefits will serve the same purpose.

VII. Capacity Payments to Suppliers: An Analogy

The concept of making upfront payments to induce participation is of course not new and a moment's reflection generates numerous examples. Magazine subscriptions come with "free gifts." Job offers to desirable candidates come with signing bonuses. Cellular phone services give a $700 phone to new customers before they have paid anything for the service. Landlords sometimes offer an initial period rent-free. All of these are simply examples to overcome inertia, motivate the expansion of a relevant choice set, transfer risk and modify the time pattern of returns, just as section 5 describes.

But from the standpoint of the management of an electric system, the obvious analogy is with capacity payments to generators. A brief examination of the history of capacity payments can help illuminate why the narrow academic view of energy-only markets fail to capture important aspects of real-world electric systems.

As mentioned above, the 1988 treatise by Schweppe *et al.*, *Spot Pricing of Electricity* formed the basis for seriously considering real-time electricity markets, as opposed to the historic practice (still the mainstay in many systems) of vertically integrated utilities which maintain control over all but the end-use decision to consume electricity. In that book, and in the elegantly formulated vision of electricity markets which followed for the next 20 years, there was little mention of payments to generators beyond payments for energy and ancillary services, including, perhaps, contemporaneous payments for capacity available which mirrored the probability of lost load on the system, or contemporaneous payments for the provision of operating reserves to give the control room flexibility in meeting its obligation to balance supply and demand.

In the actual operations of these markets, however, capacity payments, in the form of upfront promised payments to make capacity available have been implemented in all three of the Eastern US markets (PJM, NYISO and ISO-NE) and are under consideration in many more

In fact, the reasons for these capacity payments, while traceable in part to peculiarities of ISO-administered markets like price caps, are also directly responsive to the typology described above. They are largely related to bounded rationality, risk transfer and time preference, as well as a fundamental unwillingness to rely on the price system alone to make investments in time[16].

The decision to build a power plant is the decision to sink capital in a single-purpose asset. Capital seeking the highest available risk-adjusted return will look at new power plants (or the continued operation of existing ones) as but one choice in the vast set of possible choices. Power plant construction requires large resources upfront, with uncertain operating profits over a period of years. It is unsurprising that capital is unwilling to commit irreversible investments over long periods of time without an expectation of substantial net returns. The standard theorems of real options theory[17] suggest that long-lived irreversible investments require substantially higher expected returns than shorter-term or more flexible investments.

Capacity markets serve to reduce uncertainty in the prospects for new generating units by promising to transfer risk through streams of upfront payments. Unsurprisingly, they have been successful at doing so. Assessments of capacity markets which have been performed all suggest that they have increased supply of generation.

The reduction in uncertainty is by no means the only reason for capacity markets. In addition, capacity markets allow an additional stream of income not available in the energy which allows more capacity to be built than might be built by the operation of energy markets alone. This is the so-called “missing money” problem. Energy markets may not generate enough profits to justify as large a capacity margin as system planners might like. Thus, either subsidizing the entry of new capacity or, equivalently, subsidizing demand-side reductions enable planners to greatly increase the margin of generation over load.

---

[16] Preferences and the knowledge space are not particularly important on the generator side.

[17] Avinash Dixit, *Investment Under Uncertainty,* Princeton University Press, 1994

A third justification for capacity markets is to provide signals about the cost of capacity additions. These signals can in turn be used as inputs into many other decisions, particularly in transmission planning.

Note that these justifications for capacity markets would all exist even in a mandatory pricing regime. The risk of entry would still exist, the margin of supply over demand would not be freely chosen by regulatory authorities and the expected cost of new entry would still be mysterious in the absence of a capacity market even if dynamic pricing of energy were mandatory.

It is possible to run capacity markets without allowing the demand-side resources to sell into the market, but it would be inefficient and unwise. One cannot know whether the most efficient way to establish a given supply/demand balance is to increase generation or to decrease demand without taking bids from both sides.

So as long as there is a need for capacity markets which is independent of the desire for dynamic pricing, and so long as demand-side resources are a vital component of those markets, and so long as participation by those demand-side resources is voluntary, there will continue to be scope for finding out what makes customers agree to participate.

While capacity markets are an obvious source of availability payments to demand-side resources, the same underlying economic forces operate in traditionally regulated markets and government-run markets. The goal of achieving a given level of reliability at least cost cannot be solved by dynamic pricing alone.

VIII. Empirical Support

A. General results on investments in energy efficiency

The empirical support for the proposition espoused above, that availability payments are critical to securing demand response, has two fundamental components. The first is the large econometric literature devoted to the implicit discount rate for energy-saving investments. Decades of empirical research have shown that energy end-users have what seem to be very high discount rates for energy-saving investments. Most notably, Hausman (1979)[18] used observed consumer choices among air conditioning models and

[18] Hausman, Jerry. 1979. “Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables.” Bell Journal of Economics, 10 (1): 33–54.

Dubin and McFadden (1984)[19] used observed household choices among heating system to show implied discount rates of roughly 15 to 25 percent. The reasons behind this reluctance of end-users to make investments whose payback periods would seem attractive given considerably lower market interest rates have puzzled many. But, empirically, the result is quite striking.

The realization that energy end-users have significantly higher discount rates than the corporations providing them with energy led to a more-or-less standard way of implementing demand side utility programs. The utility made an investment (on behalf of ratepayers) at the customer's premises and then reaped whatever benefits were to be gained from reduced capacity expansion over time. Enhanced insulation investments, massive solar panel subsidies, smart thermostats, free energy audits, and free compact fluorescent bulbs[20] are among the host of programs that make no sense unless end users require availability payments. The primary beneficiary of any pure energy saving investment is bound to be the end user. To subsidize these investments is clear empirical proof that availability payments of some sort are required to induce socially beneficial behavior.

The failure to make investments which would appear to have favorable cost-benefit ratios is by no means limited to individual households. Commercial establishments have been demonstrated to have similar difficulties in making what would appear to be money-saving investments.

In addition, however, existing demand response programs have been much more likely to generate uptake when the upfront costs of joining are low. It is to that evidence that we turn next.

B. A Database of U.S. Demand Response Programs

The objective of our empirical analysis is to compare the rates of participation in the various types of demand response programs. To do so, we build a database using two primary sources: (1) a recent report from Navigant Research that tracks demand response programs around the world;[21] and (2) U.S. Energy Information Administration (EIA) data on U.S. retail electricity customers and providers.[22]

---

[19] Dubin, Jeffrey, and Daniel McFadden. 1984. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." Econometrica, 52 (2): 345–62.

[20] Compact fluorescents might require some subsidy even if the discount rates were low to account for inferior lighting value, but the general point still stands.

[21] Hedin M. and Woods E. Navigant Research. "Demand Response Tracker 2Q13: Global DR Programs by Country, DR Market, Resource Type, and Customer Segment." Published 2Q 2013.

[22] See U.S. EIA website: http://www.eia.gov/electricity/data.cfm

The Navigant Research report includes information on 1,263 demand response programs across the United States. The report provides the characteristics of the demand response programs, including the specific program structures (e.g. time-of-use pricing, interruptible load) and a characterization of the market in which the DR payments are made (e.g. economic or capacity markets). It also includes the entity that runs the demand response program and a characterization of this entity (e.g., investor-owned utility, cooperative, municipality). The Navigant database provides information on the U.S. state in which the program is located, the customer class at which the program is targeted, and whether the program is fully deployed or just a pilot. Finally, and most importantly for the purpose of our analysis, the Navigant database lists the number of participants in each program.

The U.S. EIA database includes the name and entity-type of each U.S. supplier of electricity and the number of retail customer for each entity by U.S. state and customer class. The database includes information on 3,114 electricity suppliers to residential customers, 3,099 electricity suppliers to commercial customers and 2,213 electricity suppliers to industrial customers**.**

To enable the comparison of participation rates across programs, we match the entities that run the demand response programs from the Navigant database with the entities that supply electricity from the EIA database. We then estimate a participation rate for each demand response program by dividing the number of participants by the total customers in the relevant U.S. state and customer class.[23]

Various demand response programs are removed from the Navigant database. These include all "pilot" programs, all programs for the "agricultural" customer class, and programs for which no data were found on number of participants or the total customer base.

We are left with 943 demand response programs, and information on the characteristics of each program, including the estimated participation rate. In the following section we describe the results of our analysis of this database of demand response programs. [24]

C. Analysis of Participation in Demand Response Programs

Our analysis of U.S. demand response programs shows that programs with availability payments have substantially higher take rates than programs with "economic market" payments. This result is directly in line with the hypothesis that availability payments are

---

[23] Note that many demand response programs have specific eligibility restrictions for participation in demand response programs. For example, only customers that use above a certain amount of electricity may be eligible for certain programs. It was not feasible to account for these specific requirements. To the extent that these restrictions limit the number of customers that could have participated in demand response programs, our estimated participation rates will be too low.

[24] Details on the data procedure are given in the Appendix.

required for participation, because programs with availability payments garner expected benefits before the costs are incurred.

In our sample, total participation in capacity programs is 4.8 million customers, compared to 1.5 million customers for economic programs, as displayed in Figure 1.[25]

Figure 1: Total Participation in Demand Response Programs

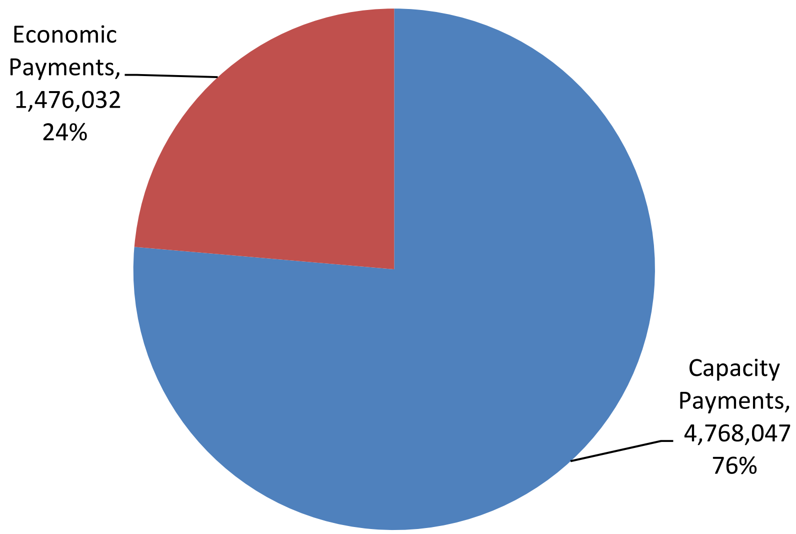

Figure 2 shows that overall participation rates are estimated at 7.1% for capacity programs compared to 1.9% for economic programs. The average program participation rates are 11.2% for capacity programs and 2.5% for economic programs.

[25] As noted above, we estimate "participation rates" by dividing total customer participation in demand response programs by the total customers of the demand response program providers (in the customer class at which the program is directed).

Figure 2: Participation Rates in Demand Response Programs

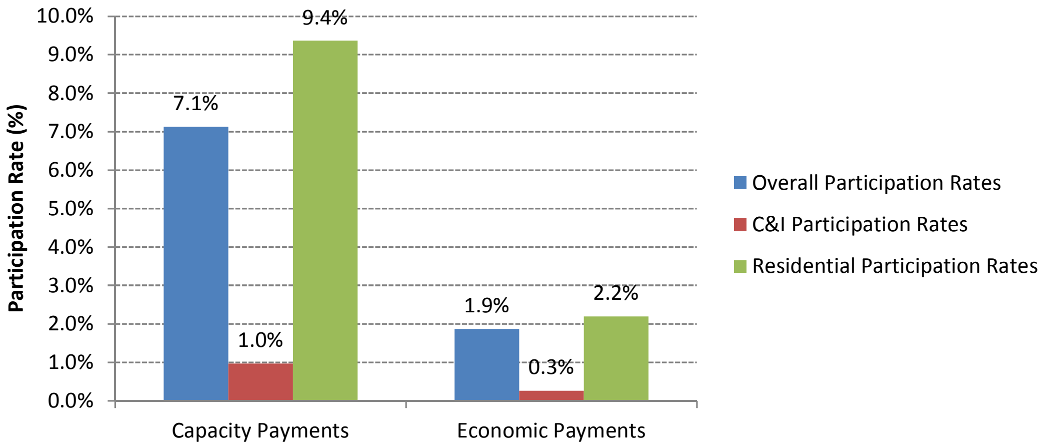

Note that this outcome is not because economic programs are rare; in fact, Figure 3 shows that economic programs cover a higher number of potential customers than capacity programs. Both programs are in fact quite common, so the fact that capacity programs do much better is almost surely the result of behavioral issues that favor capacity based programs over economic programs.

Figure 3: Total Potential Participation in Demand Response Programs

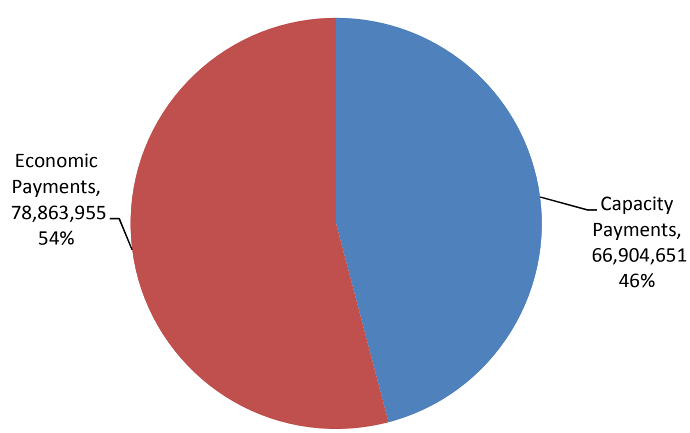

When the programs are divided by customer class, we find much higher participation rates for programs directed at residential customers compared to commercial and industrial customers, and that the residential programs are responsible for the majority of the difference in participation rates between capacity programs and economic programs—Figure 2 shows participation rates are 9.4% for residential capacity programs versus 2.2% for residential economic programs. Overall participation rates for programs directed at commercial and industrial customers are closer in absolute terms, though still over three times higher for capacity programs (1.0%) versus economic programs (0.3%).

The larger discrepancy in participation rates for residential customers again supports the hypothesis that availability payments are required for participation. Providers of residential programs have always understood that behavioral uptake rates require upfront incentives. By contrast, commercial and industrial programs have been assumed to be more "bottom-line" focused, requiring lower levels of incentives on the grounds that a commercial or industrial establishment could be presumed to seek out savings on its own. The take rates shown here indicate that the case for that proposition is weak.

When the participation rates of the programs with availability payments are further parsed, Table 1 shows that direct load control programs have much higher overall participation rates (9.2%) than interruptible load programs (0.2%). Direct load control programs are among the most likely to make upfront investments on the consumer's behalf. With regard to the demand response provider, we find that overall participation rates are significantly larger when the provider is a cooperative (15.4%) compared to a municipality (7.8%) or an investor-owned utility (6.6%). Breaking out the programs by region, the South has somewhat higher participation rates in capacity programs (9.7%) compared to the Northeast (7.4%), the Midwest (7.2%), or the West (4.4%).

Table 1: Demand Response Programs with Availability Payments

	Total Participation Rate	Average Participation Rate	Number of Programs
Total Participation Rate:	7.1%	11.2%	503
Participation Rates by Program Type:			
Direct Load Control	9.1%	16.1%	310
Interruptible Load	0.2%	3.5%	167
Load as Capacity	1.2%	3.1%	22
Participation Rates by Provider Type:			
Investor-Owned Utility	6.6%	4.3%	200
Cooperative	15.4%	14.8%	211
Municipality	7.8%	18.8%	69
Participation Rates by Region:			
Northeast	7.4%	9.6%	47
Midwest	7.2%	13.9%	282
South	9.7%	7.6%	127
West	4.4%	5.9%	44
Participation Rates by Customer Class:			
Residential	9.4%	17.8%	235
Commercial & Industrial	1.0%	5.5%	268

Notes:

Total participation rates are calculated as the sum across all programs of total program participants divided by the sum across all programs of total potential participants (i.e. customers in the relevant customer class).

Average participation rates are calculated by dividing program participants by total potential participants for each individual program, and then averaging the results across all programs.

A more detailed examination of the participation rates of the economic programs is shown in Table 2. Time-of-use programs are more widespread and have higher overall participation rates (2.1%) than real-time pricing (0.2%), critical-peak pricing (0.3%) or demand bidding and buy-back programs (0.2%). To some extent of course, the paucity of take rates for real-time pricing is no doubt related to the relative young age of these programs. Nevertheless, there is no evidence that general acceptance of real-time pricing is anywhere close to reality.

Note that we do not find significant differences in participation rates for economic programs across different types of demand response providers. The overall participation rates of economic demand response programs in the West (6.1%) and Northeast regions (5.6%) are higher than those in the South (0.3%) or Midwest (0.7%) regions.

Table 2: Demand Response Programs with Economic Payments

	Total Participation Rate	Average Participation Rate	Number of Programs
Total Participation Rate:	1.9%	2.5%	436
Participation Rates by Program Type:			
Time-of-use Pricing	2.1%	2.9%	364
Critical-Peak Pricing	0.3%	0.9%	43
Real-time Pricing	0.2%	0.2%	24
Participation Rates by Provider Type:			
Investor-Owned Utility	1.7%	1.8%	221
Cooperative	1.7%	3.2%	106
Municipality	1.0%	2.1%	91
Participation Rates by Region:			
Northeast	5.6%	3.6%	45
Midwest	0.7%	1.6%	191
South	0.3%	2.8%	134
West	6.1%	4.4%	58
Participation Rates by Customer Class:			
Residential	2.1%	3.4%	151
Commercial & Industrial	0.9%	2.1%	284

Notes:

Total participation rates are calculated as the sum across all programs of total program participants divided by the sum across all programs of total potential participants (i.e. customers in the relevant customer class).

Average participation rates are calculated by dividing program participants by total potential participants for each individual program, and then averaging the results across all programs.

IX. Conclusion

Availability payments, whether directly in the form of monetary payments, or indirectly in the form of technology or technological assistance, are currently the surest path to raising participation in a demand-side program. Without such incentives, the hypothetical reliance on self-interest is unable to overcome well-understood behavioral barriers to change. This fact has been the case from the beginnings of demand-side programs and continues to be true today.

Availability payments are needed wherever the control room would like the ability to use voluntary controls on demand and not simply rely on additional generation to keep supply and demand in balance. This is really not a question of industry structure. It flows from the physical requirement to keep supply and demand in balance. This physical requirement is present in every electric system, whether it uses an independent system operator (ISO), a traditional vertically integrated structure, or a pure public system. Voluntary reductions of demand can save money at the margin of course, which can make a much larger cost difference in systems in which customers are charged a price which reflects marginal cost, but the benefit of demand-side participation to real system reliability is critical to all electric systems, simply because the social costs of involuntary load shedding are so high.

The siren call to real-time pricing by academic economists does not eliminate the need for such programs, at least not those which provide the most valuable commodity for the control room: dispatchable capacity. Dynamic pricing has its own potential benefits, but the real-time reliability benefits of dispatchable demand programs cannot be found in these programs. Further, at least for the moment, voluntary dynamic pricing implementation has been slow and even its most strenuous advocates are not optimistic about rapid growth. Voluntary demand reductions supported by availability payments will remain for the foreseeable future a vital part of the toolkit of any control room which seeks to keep supply and demand balanced.

Appendix: Methodology for Demand Response Participation Rate Analysis

This appendix provides detailed information on the methodology used to conduct our empirical analysis on the participation rates of U.S. demand response programs.

A. Data Sources

Our analysis uses two primary sources:

1. A 2013 report from Navigant Research entitled "Demand Response Tracker 2Q13: Global DR Programs by Country, DR Market, Resource Type, and Customer Segment." [26] This report tracks demand response programs around the world, including a spread-sheet of information on 1,263 demand response programs across the United States. The report provides the characteristics of the demand response programs, including the specific program structures (e.g. time-of-use pricing, interruptible load) and a characterization of the market in which the DR payments are made (e.g. economic or capacity markets). It also includes the entity that runs the demand response program and a characterization of this entity (e.g., investor-owned utility, cooperative, municipality). It provides information on the U.S. state in which the program is located, the customer class at which the program is targeted, and whether the program is fully deployed or just a pilot. Finally, and most importantly for the purpose of our analysis, the Navigant database lists the number of participants in each program.

2. U.S. Energy Information Administration (EIA) data on U.S. retail electricity customers and providers.[27] The EIA database includes the name and entity-type of each U.S. supplier of electricity and the number of retail customer for each entity by U.S. state and customer class. The database includes information on 3,114 electricity suppliers to residential customers, 3,099 electricity suppliers to commercial customers and 2,213 electricity suppliers to industrial customers**.**

B. Data Cleaning

To enable the comparison of participation rates across programs, we match the entities that run the demand response programs from the Navigant database with the entities that supply electricity in the same U.S. state from the EIA database.

---

[26] Hedin M. and Woods E. Navigant Research. "Demand Response Tracker 2Q13: Global DR Programs by Country, DR Market, Resource Type, and Customer Segment." Published 2Q 2013.

[27] See U.S. EIA website: http://www.eia.gov/electricity/data.cfm

However, this left a significant number of problems which had to be manually corrected. In various instances the attempt to match demand response program providers with electricity suppliers returned errors. In many cases, the error was simple. For example, there were several cases in which one source used the abbreviation “Coop” in the entity name and the other used the full word “Cooperative.”

After this process was completed, we examined the remaining unmatched records a second time. These records were matched with customer data from a secondary EIA source, the EIA-861 Annual Electric Power Industry Reports.[28] For example, the information on customers for “Bangor Hydro Electric Company” was found only in this second EIA source.

Unmatched records remained even after the second-round examination. These records were concentrated among the following four larger utilities and cooperatives: Xcel Energy, Columbus Southern Energy Company, Dairyland Power Cooperative, and WPPI Energy.

Dairyland Power Coop and WPPI Energy post lists of member municipalities on their webpages. These were matched directly with EIA customer data and summed to determine potential participation, in the same manner described above.

For Xcel Energy, customer data was available only in the form of total customers for each individual state in Xcel’s service territory. We approximated the residential and commercial & industrial customer split using the average splits in the same states but for different utilities. These averages were then applied to the data on Xcel’s total customer base.

The final step is to estimate a participation rate for each demand response program by dividing the number of participants by the total customers in the relevant U.S. state and customer class.

Note that many demand response programs have specific eligibility restrictions for participation in demand response programs. For example, only customers that use above a certain amount of electricity may be eligible for certain programs. It was not feasible to account for these specific requirements. To the extent that these restrictions limit the number of customers that could have participated in demand response programs, our estimated participation rates will be too low.

---

[28] See U.S. EIA website: http://www.eia.gov/electricity/data/eia861/

C. Summary of the Data Cleaning Process

The following provides a summary of the data that were eliminated from our dataset before we conducted our statistical analysis:

- All demand response programs characterized as “pilot” programs (29 in total);
- All programs for which the customer segment was characterized as “Agriculture” (61 in total);
- All programs with no data on number of participants (210 in total);
- All programs for which entity names could not be matched between the Navigant and EIA sources (154 in total);
- All programs for which the estimated participation rate is over 100% (13 in total).

In all, 320 records were deleted and 943 records remained. Please note that the categories listed above are not necessarily mutually exclusive. Therefore, the list double-counts some records and the record count sums to more than 320.

NERA Economic Consulting
1166 Avenue of the Americas
New York, New York 10036
Tel: 1 (212) 345-3000 Fax: 1 (212) 345-4650
www.nera.com