MARKETS AND PRICING FOR INTERRUPTIBLE ELECTRIC POWER

by

Thomas W. Gedra and Pravin P. Varaiya

Memorandum No. UCB/ERL/IGCT M91/11

12 February 1991

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

MARKETS AND PRICING FOR INTERRUPTIBLE ELECTRIC POWER

by

Thomas W. Gedra and Pravin P. Varaiya

Memorandum No. UCB/ERL/IGCT M91/11

12 February 1991

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

Markets and Pricing for Interruptible Electric Power

Thomas W. Gedra Pravin P. Varaiya
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley

Abstract

We propose a market for interruptible, or callable, forward contracts for electric power, in which the consumer grants the power supplier the right to interrupt a given unit of load in return for a price discount. The callable forward contracts are traded continuously until the time of use. This allows recourse for those customers with uncertain demand, while risk-averse consumers can minimize their price risk by purchasing early. Callable forward contracts are easily enforced, and can be directly incorporated into the utility's economic dispatch procedure.

1 Introduction

In existing power systems, the required generation capacity is usually much larger than the average amount of power actually provided. This is due to two factors: load variation over time (or load factor) and the need for a nonzero "reserve margin," or spare capacity, to be used in the event of unanticipated contingencies (unexpectedly high demand or failure of generation components). These are important because electric power is storable only at very high cost, if at all, requiring production to equal consumption on a second-to-second basis. Thus, installed generating capacity must always exceed the annual peak demand (plus a reserve margin), even though some of the capacity is idle for much of the time.

If it is uneconomic or infeasible to build sufficient capacity, occasions will arise in which demand exceeds available supply. It is then necessary to remove some of the load demand to maintain system security. Methods currently in use include voltage reduction and rotating outages. Both methods are fairly indiscriminate, and result in welfare losses. If rationing occurs infrequently, these welfare losses are relatively small. However, if generation capacity lags behind demand, service interruptions will become more frequent, and so the need to efficiently ration electric power will become more important.

Submitted to the IEEE Power Engineering Society Summer Meeting 1991.

A number of proposals have been made to reduce system load in a socially efficient manner. We will refer to these load-reduction schemes as "demand-side management" (DSM) [1]. Perhaps the simplest DSM technique is conservation through efficiency improvements (in lighting, motors, etc.). Other examples of DSM include direct load control and cycling of certain loads.

Other DSM techniques may be described as price-based schemes, because they provide the consumer with monetary incentives to help reduce system peaks. The simplest scheme, at least conceptually, is spot pricing: through the operation of a "spot" market, price would instantaneously adjust until the market "clears" (*i.e.* demand equals supply) [2]. Under certain standard assumptions, it can be shown that spot pricing achieves a social-welfare maximizing allocation.

A variety of other schemes have been proposed to supplement or replace spot pricing. The use of forward contracts has been suggested to allow users to schedule their consumption in advance and to ease some of the price risk inherent in spot pricing [3]. The use of priority service contracts, in which the consumer selects the probability of service interruption according to the need for reliable power (and pays accordingly), has also been proposed [4]. The pricing method discussed here is closely related both to forward contracts and priority pricing, but has certain implementational advantages, such as ease of enforcement and the ability to be incorporated into economic dispatch.

2 Callable Forwards

The basic market instrument we propose is a type of interruptible-load contract, which we call a *callable forward*. A callable forward is a bundle of two contracts. The first of these is a *forward contract*, which is owned by the consumer, and which guarantees that the utility will deliver to the consumer one unit of electrical energy at a particular time T, the *delivery date*. The price paid for this unit is determined at the time the contract is made. The second contract is a *call* on the same unit of energy. A call confers the right, but not the obligation, to purchase a given commodity at a given price, called the *strike price*, which we denote by k. The call portion of a callable forward is sold by the consumer back to the utility; we describe the consumer's resulting obligation by saying that the consumer is *short* a call. Using this terminology, a callable forward

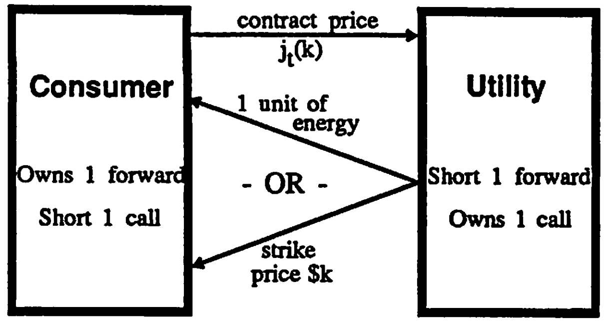

Figure 1: Contractual obligations for a callable forward.

is a bundle of a forward and a short call.

Thus a consumer who owns a callable forward is guaranteed to receive from the utility at time T, either one unit of energy or $\$k$ (and no energy), at the option of the utility. In the second case, we say that the utility has exercised its call option at a cost of $\$k$, thus relieving itself of the obligation to provide energy. If the utility does not exercise the option, it remains obligated (by the forward contract) to deliver the energy. These contractual obligations are shown in Figure 1.

Conceptually it may be helpful to envision a centralized implementation of callable forward contracts, in which there is a single utility or price-setting authority, which we will call the "price-setter." The price-setter unilaterally sets the prices of the callable forwards based on its best information about the future spot price. Consumers respond by purchasing the appropriate collection of callable forward contracts, taking the prices as given. To make this decision, the consumer trades off the probability of interruption (which is lower for higher k) against the cost of the contract (which is higher for higher k). Consumers will then select contracts as we describe below: consumers with a lower cost of interruption will choose contracts with a lower k, and so will be interrupted first in the event of a shortage, since they are less expensive for the utility to interrupt. This results in socially optimal rationing.

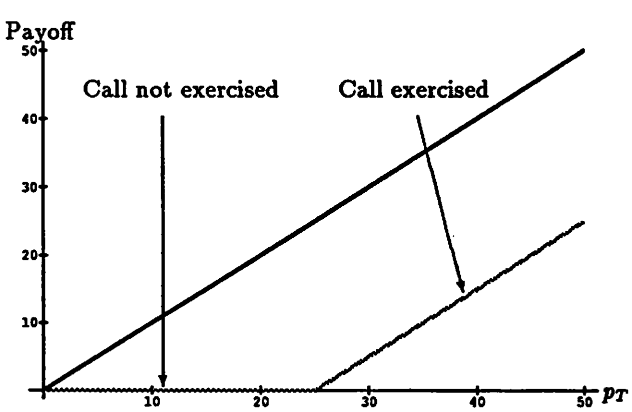

Figure 2: Payoff of forward (black) and call (gray, k=\$25/MWh) vs. spot price p_T (Both in \$/MWh).

3 Pricing callable forwards

A common device used to depict contracts and options is the *payoff diagram* [5]. A payoff diagram is simply a graph of the value, or price, of the contract or option as a function of the price of energy at time T. For such a diagram to make sense in our context, we must assume that there is a well-defined price of a unit of electric energy at time T, which will be denoted by p_T. This price could be a "spot" price or shadow price. Figure 2 shows the payoff diagram for a forward contract and for a call option with strike price \$25/MWh written on a unit of electrical energy deliverable at time T.

At the time of delivery T, a forward contract must have a price exactly equal to the spot price of the underlying commodity. Thus if f_t denotes the price of the forward contract at time t, we must have

$$f_T = p_T, \tag{1}$$

so that the payoff diagram for the forward is just a line through the origin with unit slope.

To evaluate the payoff of a call option at time T, we must consider two cases. In the first case, the spot price p_T is greater than the strike price k of the call. Exercising the call in this case yields a unit of energy for only k, which could be sold on the open market for p_T. Thus, the call is worth $p_T - k$. In the second case, the spot price p_T is less than or equal to the strike price k of the call. Since a unit of electrical energy may be purchased on the open market for less than the strike price of the call, the call is not exercised and becomes worthless. (Note that nothing prevents the holder of the call from exercising it even in this second case, but this would be foolish, as it would result in a loss of $k - p_T$.) To summarize these two cases, we can write

$$c_T(k) = \max\{0, p_T - k\}, \tag{2}$$

where $c_t(k)$ denotes the price of a call having strike price k at time t.

To find the payoff diagram of a callable forward with strike price k, we just subtract the payoff of the call from

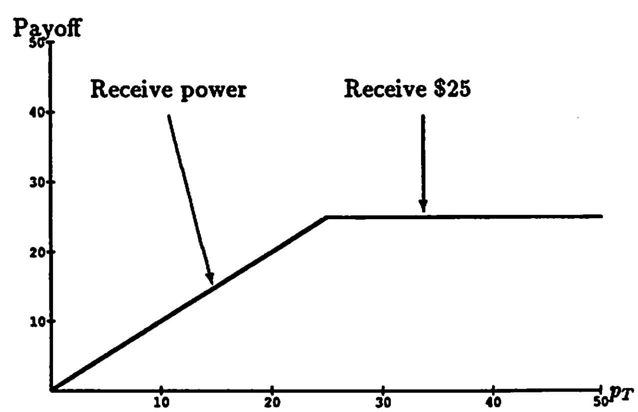

Figure 3: Payoff for callable forward (k=\$25/MWh) vs. spot price p_T (Both in \$/MWh).

that of the forward. The resulting payoff can be seen in Figure 3. Mathematically, we have

$$j_T(k) = f_T - c_T(k) = p_T - \max\{0, p_T - k\} = \min\{p_T, k\}, \tag{3}$$

where $j_t(k)$ denotes the price of a callable forward with strike price k at time t.

We have discussed the time T payoff of a callable forward contract. We now examine the price of such a contract when it is purchased at some time $t < T$.

First, define the conditional distribution of the terminal spot price by

$$Q_t(k) = \text{Prob}\{p_T \le k | \mathcal{H}_t\}, \tag{4}$$

where $\mathcal{H}_t$ denotes all that is known at time t, and define the associated density $q_t(k) = \frac{\partial}{\partial k}Q_t(k)$, which we assume exists.

We assume that consumers price a particular contract at time t as the expected value of the terminal (time T) payoff of the contract, conditional on what is known at time t. So the price of a forward contract is

$$f_t = \mathrm{E}[p_T|\mathcal{H}_t] = \int_0^\infty p\, q_t(p)\, dp. \tag{5}$$

Similarly, for the callable forward contract we obtain

$$j_t(k) = \mathrm{E}[j_T(k)|\mathcal{H}_t] = \int_0^\infty \min\{p, k\}\, q_t(p)\, dp. \tag{6}$$

Just as knowledge of the conditional density $q_t(k)$ allows calculation of the price of a callable forward contract $j_t(k)$, knowledge of contract prices can be used to infer the density $q_t(k)$. To see this, we first integrate (6) by parts to yield the expression

$$j_t(k) = \int_0^k [1 - Q_t(p)]\, dp. \tag{7}$$

Differentiation of (7) yields

$$\frac{\partial}{\partial k} j_t(k) = 1 - Q_t(k), \tag{8}$$

and, differentiating again,

$$q_t(k) = -\frac{\partial^2}{\partial k^2} j_t(k). \tag{9}$$

Additional properties of $j_t(k)$ which follow from its definition and the discussion above are:

- $j_t(k)$ is nondecreasing and concave in k.
- $j_t(k) \le k$, $\forall k$.
- $\lim_{k\to\infty} j_t(k) = f_t$.

Figure 4 illustrates how the price curve for callable forwards varies with the uncertainty of the spot price p_T. The lighter gray curves are associated with more uncertainty than the darker curves; as the uncertainty decreases, the price curves approach the kinked graph of $\min\{f_t, k\}$. Since the uncertainty of p_T generally decreases as $t \to T$, Figure 4 may also be interpreted as a possible time evolution of the curves $j_t(k)$, with darker curves representing later times.

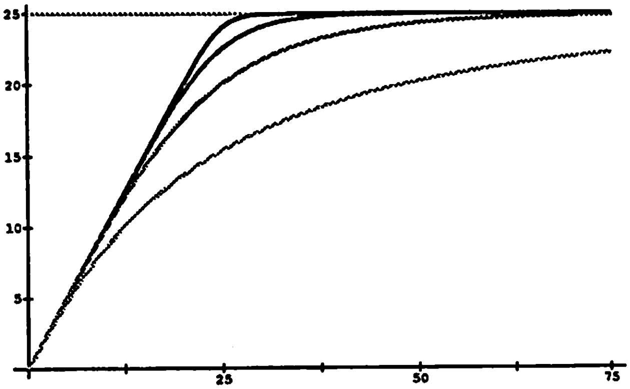

Figure 4: Price vs. k for different amounts of uncertainty ($f_t = \$25$).

4 Consumer Self-Selection

A desirable feature of any market-based allocation scheme is that individuals' choice of contracts collectively maximizes social welfare. That is, each consumer freely chooses the same contract that he or she would be assigned in a social-welfare maximizing allocation. For a consumer who has a single unit of load having valuation $\$v$, spot pricing, or any other optimal allocation, will allocate a unit of energy whenever $p_T < v$, and will not allocate the unit when $p_T > v$. Since a callable forward will be called when $p_T > k$, a callable forward with $k = v$ is the contract which yields this allocation. Thus, we would like the consumer to make the selection $k = v$.

We show this is the case if the consumer's demand is independent of p_T, under the assumption that the density $q_t(k) > 0$ whenever $k > 0$. First we examine the case in which a consumer has only a single unit of load, which will yield a benefit or value to the consumer of $\$v$ if the load is supplied, and which yields a benefit of $0 if is not supplied. If such a consumer purchases a unit callable forward having strike price k, the resulting monetary benefit will be

$$\text{Benefit} = B_t(k) = \begin{cases} v - j_t(k) & \text{if service is received,} \\ k - j_t(k) & \text{otherwise.} \end{cases} \tag{10}$$

Because the consumer's contract will be called if $p_T > k$, the consumer will receive power with probability $Q_t(k)$. If the consumer is risk-neutral (*i.e.*, the consumer's perceived benefit equals the expected monetary benefit), then he or she will choose a contract which maximizes

$$\mathrm{E}[B_t(k)|\mathcal{H}_t] = k - j_t(k) + (v - k)Q_t(k). \tag{11}$$

A necessary condition for an interior solution to this maximization is that

$$\begin{aligned} 0 &= \frac{\partial}{\partial k}\mathrm{E}[B_t(k)|\mathcal{H}_t] \\ &= 1 - \frac{\partial}{\partial k} j_t(k) - Q_t(k) + (v - k)q_t(k). \end{aligned} \tag{12}$$

Recalling from (8) that $\frac{\partial}{\partial k} j_t(k) = 1 - Q_t(k)$, this becomes

$$0 = (v - k)q_t(k). \tag{13}$$

Using the assumption that $q_t(k) > 0$, $\forall k > 0$, we obtain

$$k = v. \tag{14}$$

Further, since $\frac{\partial}{\partial k}\mathbf{E}[B_t(k)|\mathcal{H}_t] = (v-k)q_t(k)$, it follows that

$$\begin{aligned} \frac{\partial}{\partial k}\mathbf{E}[B_t(k)|\mathcal{H}_t] &> 0 \text{ if } k < v \\ &< 0 \text{ if } k > v. \end{aligned} \tag{15}$$

Finally, if the consumer chooses $k = v$, the resulting utility will be

$$\mathbf{E}[B_t(v)|\mathcal{H}_t] = v - j_t(v), \tag{16}$$

which is positive for $v > 0$ by virtue of the assumption that $q_t(k) > 0$, $\forall k > 0$, so that the consumer with $v > 0$ will always choose to purchase the contract with $k = v$ rather than purchase no contract. Thus we have shown that the choice $k = v$ is the unique maximizer of (11). So, in this case, the consumer self-selects the optimal contract.

It is worth noting that the selection $k = v$ is a riskless contract for the consumer, since he or she receives either electricity (with a utility of \$$v$) or \$$k$ = \$$v$ in cash. Since even risk-neutral consumers choose a riskless allocation, one might suspect that risk-averse consumers would make the same choice. This is, in fact, the case, as is shown in Appendix A.

Until now we have assumed that the consumer's valuation v was a deterministic quantity. We now consider the case where v is a random variable, which may or may not be independent of p_T. We again ask which contract the consumer will select.

In the case where v is independent of p_T the self-selection is quite similar to that in (14), except that v is replaced by its conditional expectation. Thus we have

$$k = \mathbf{E}[v|\mathcal{H}_t], \tag{17}$$

which may be interpreted as the consumer's "best guess" of the true value of the unit. This result is a special case of (32) in Appendix B. It can be shown, using the same argument as above, that (17) is, in fact, the unique maximizer of the consumer's benefit.

Uncertainty in v which is uncorrelated with p_T may result from factors which are unique to a particular consumer (provided that the consumer is small relative to the total system demand). For example, an industrial user may not know whether its machinery will fail before the time of use, or it may not be known whether a shipment of raw materials for a manufacturing process will arrive in time for a production run. If the machinery fails, or raw materials are not available, the benefit derived from a unit of electrical energy will obviously be reduced. Correlated uncertainty may result either from the consumer being large enough to affect the market price, or from systemwide factors which affect a significant portion of the system load, such as ambient temperature.

In the general case, where v and p_T are not independent, the necessary condition for an interior maximum of (29) is

$$k = \mathbf{E}[v|\mathcal{H}_t, p_T = k]. \tag{18}$$

This result is derived in Appendix B. Because the right hand side of (18) depends on k, there may be no solutions

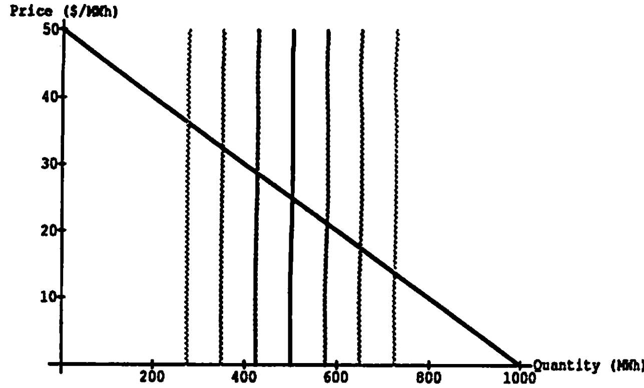

Figure 5: Demand and supply curves for section 5; price is in \$/MWh, quantity is in MWh.

for k, or there may be multiple solutions. If there is a solution, it may not be the global maximizer of the consumer's benefit. Appendix B also gives an example to show that the consumer may "hedge", by buying a contract with higher or lower k, depending on whether the correlation with p_T is positive or negative, respectively. When this occurs, the consumer will receive a suboptimal allocation (*i.e.* at time T, we will not have $k = v$).

5 An Example

To illustrate our results, we now turn to an example. Let T=1. We assume that the aggregate demand curve is deterministic, and is given by

$$p(q) = \frac{1}{20}[1000 - q] = 50 - \frac{q}{20}, \tag{19}$$

where $p(q)$ is the price (in \$/MWh) corresponding to a given quantity q (in MWh). This is the demand curve corresponding 1000 MWh of load, with values uniformly distributed between 0 and \$50 per MWh. On the supply side, we assume, for simplicity, that electricity has zero variable cost, but that the generation capacity available is limited to some value s. An equivalent assumption, at least for calculating prices, is that variable cost is nonzero, but that the capacity constraint is always binding, so that the demand curve always intersects the vertical portion of the supply curve. We assume that s is a Gaussian random variable, with mean 500 MWh and standard deviation 100 MWh. The supply and demand curves are illustrated in Figure 5. We assume that s is actually the terminal value s_1 of a rescaled Brownian motion [6], given by

$$\begin{aligned} ds_t &= 150\, dW_t, \text{ and} \\ s_0 &= 500, \end{aligned} \tag{20}$$

where W_t denotes standard Brownian motion. In terms of s, the spot price is simply

$$p_T = p_1 = p(s_1) = 50 - \frac{s_1}{20}. \tag{21}$$

Thus we can write the forward price as

$$f_t = \mathbf{E}[p_T|\mathcal{H}_t] = 50 - \frac{s_t}{20}. \tag{22}$$

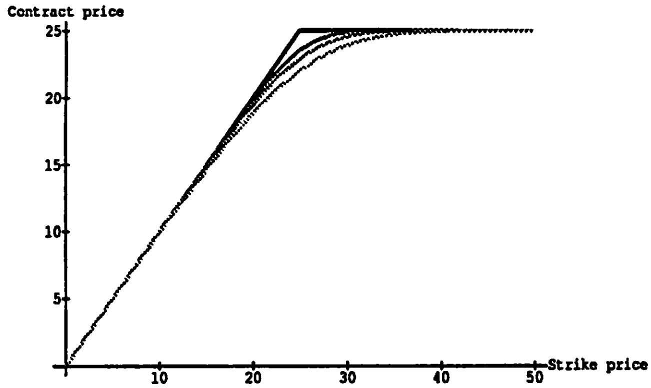

Figure 6: Callable forward prices vs. strike price k at t=0,1/2,3/4 and 1.

Because s is Gaussian and $p(\cdot)$ is linear, we can write the conditional density of p_T given $\mathcal{H}_t$ as

$$q_t(p) = \frac{1}{7.5\sqrt{2\pi(1-t)}}\exp[-\frac{1}{2}\frac{(p-f_t)^2}{(7.5)^2(1-t)}]; \qquad (23)$$

that is, p_T, given $\mathcal{H}_t$, has a Gaussian distribution with mean f_t and variance $(7.5)^2(1-t)$.

We can now compute $j_t(k)$ from (23) by integrating once to obtain $Q_t(k)$, then integrating $1-Q_t(k)$ to obtain $j_t(k)$. Figure 6 shows numerically obtained plots of $j_t(k)$ as a function of the strike price k for $t = 0, \frac{1}{2}, \frac{3}{4}$, and 1, with price curves for larger t shown in darker gray.

6 Contract Repurchasing

So far, we have considered the decision of the consumer who first purchases a contract at time t. However, it may be the case that the consumer already holds a contract at time t. We now examine whether the old contract should be held or a new one purchased at time t. Suppose the consumer's original contract was purchased at some time $t_0 < t$, and that the strike price of the original contract was k_0. It is assumed that the old contract may be sold back to the market at the current price. The consumer's net benefit from the two transactions will be

$$\begin{aligned} B_t(k) &= k - j_t(k) + (v-k)1\{p_T < k\} \\ &\quad + [j_t(k_0) - j_{t_0}(k_0)], \end{aligned} \qquad (24)$$

where k is the strike price of the new contract purchased at time t. The bracketed term above represents the gain or loss resulting from holding contract k_0 from time t_0 to time t; notice this term does not depend on k. Thus, the gain or loss term above is "sunk" (the consumer has already incurred it, and cannot affect it by the choice of the new k). The consumer maximizing $\mathrm{E}[B_t(k)|\mathcal{H}_t]$ will thus choose the same k as in the case where no previous contract existed. Note that the maximization of $\mathrm{E}[B_t(k)|\mathcal{H}_t]$ actually includes two other options for the consumer: choosing $k = k_0$ corresponds to the decision not to change contracts, while choosing $k = 0$ corresponds to selling one's contract without repurchasing (since a contract with $k = 0$ is guaranteed to be interrupted, and is thus equivalent to no contract at all).

For consumers with deterministic v, this implies that, once the original contract purchase is made, the consumer will hold the contract until the time of use. Such consumers initially purchase the optimal contract, and never change contracts, regardless of price changes. On the other hand, consumers with nondeterministic demand which is independent of p_T will trade continuously, once they have initially purchased a contract, in such a way that at each time t they hold a contract with $k = \mathrm{E}[v|\mathcal{H}_t]$. If we assume that information is revealed continuously (no shocks), so that $\lim_{t\to T}\mathrm{E}[v|\mathcal{H}_t] = v$, such consumers ultimately hold the optimal contract (*i.e.*, $k = v$) at time T.

7 Implementation

There are two basic mechanisms for implementing the contracts described above: a centralized implementation and a market implementation.

In the centralized implementation, mentioned earlier, we assume that there is a single price-setter which unilaterally sets the price curve $j_t(k)$ based on its best information about the future spot price, summarized by $Q_t(k)$. Consumers respond by self-selecting and purchasing the appropriate collection of callable forward contracts, taking the price curve $j_t(k)$ as given. If consumers believe that the price-setter's basis for setting prices, namely $Q_t(k)$, is correct, they will self-select as described above. In particular, consumers whose demand is independent of future spot price will truthfully reveal their best estimate of their demand.

A major drawback of the centralized implementation is that consumers may themselves have information concerning the future spot price. While some of this information may come from the price-setter's $Q_t(k)$, the individual may possess some information which the price-setter does not. In this case, the consumer's $Q_t(k)$ will be different from that of the price-setter; consequently the consumer's $Q_t(k)$ will not be consistent with the price curve $j_t(k)$ set by the price-setter, in the sense that (8) will no longer relate the consumer's $Q_t(k)$ and the price-setter's $j_t(k)$. There is no guarantee, under these circumstances, that consumer self-selection will result in a socially optimal allocation of electrical energy.

In a market implementation, there may be many utilities or energy suppliers, without the necessity of a price-setting authority. The producers and consumers interact via a market, in which any two participants are free to trade at a price and time of their own choosing. This market need not, of course, be a physical market like the New York Stock Exchange – it could be an electronic market, in which bid and ask quotes are exchanged electronically. Each market participant has his or her own information regarding the future spot price, and thus possesses an individual $Q_t(k)$. Each participant thus also possesses a price curve $j_t(k)$, related to his or her own $Q_t(k)$ through (8). As trading occurs, information will be exchanged through the market, causing participants to update their $Q_t(k)$, and hence $j_t(k)$, accordingly. In a perfect market, every participant has the same $Q_t(k)$ and the same $j_t(k)$ as every-one else, as a result of this information exchange, and this becomes public information. The resulting market $Q_t(k)$ thus reflects a "consensus" view, reflecting the information possessed by all market participants.

Regardless of which implementation is used, we assume that trading is permitted on a continuous basis from the time the contracts are initially offered (time 0) until the time of use (time T). While many participants may simply purchase a contract which is held until time T, some consumers will choose to dynamically adjust their holdings. The ability to do this provides recourse to those consumers who find they have chosen the "wrong" contract, due to lack of information at the time of contract purchase.

Callable forward contracts possess the useful property that they can be directly implemented and easily enforced. Since the seller of the contract must either provide power or pay the strike price k to the contract holder, it is easy to determine whether the contract has been breached. This is in contrast to interruptible-load contracts which are, in theory, specified in terms of the probability of receiving service, such as priority service contracts [4] or the interruptible-load contracts described in [7]. Since it is difficult to determine whether a contract specified in terms of probabilities has been breached, such contracts are typically implemented by specifying the frequency of interruptions, total number of interruptions, or total duration of interruptions. However, such contractual approximations are not directly addressed by the underlying theory [1].

Finally, an extremely important property of callable forward contracts is that they may be integrated into economic dispatch. In economic dispatch, when an increment of power is required, it is supplied by the generating unit with the lowest incremental cost. A callable forward contract, as a dispatchable demand-side resource, has a well-defined incremental cost, namely the strike price k. Thus, dispatch of callable forwards by the utility could be easily integrated into the usual economic dispatch procedure, even though it is a demand-side resource.

8 Conclusions

We have proposed a type of interruptible-load contract, the callable forward, which has a number of useful properties. It is a contract which can be directly implemented and easily enforced. The decision to call (interrupt) a given set of callable forwards is easily made within utilities' existing economic dispatch procedures. By permitting continuous trading in callable forwards, consumers have recourse in the event that their initial contract selection is no longer appropriate.

For consumers whose future demand is independent of the future spot price, self selection yields a contract which provides power to a consumer if and only if the spot price is less than that consumer's valuation of that power. This allocation is identical to that achieved under spot pricing, which is known to yield a socially optimal allocation. In spite of the fact that there is a risk of power interruption, these consumers suffer no financial risk as a result of their contract purchase (since they choose $k = v$); they choose $k = v$ even if they are risk-averse.

Self-selection conditions were also derived for consumers whose demand is correlated with the spot price, but in general these selections do not result in selection of the "correct" contract: consumers choose k greater or less than v, depending on whether the correlation between their demand and the spot price is positive or negative.

Future work will address the question of what other options or financial instruments may need to be introduced into the market in order to induce these consumers to self-select in a socially optimal way. We also plan to examine what type of options or instruments would be appropriate for decision-makers on the supply side of the power system, such as cogenerators of independent power producers, focusing particularly on the issue of dispatchability.

Appendix A: Risk-Averse Consumers

As before, we assume that the consumer has a single unit of load with valuation v. We assume that the consumer's attitude towards risk is that he or she chooses a contract which maximizes the expected utility of the outcome, where the utility function $U(\cdot)$ is a strictly increasing, concave function, having derivative $u(x) = \frac{\partial}{\partial x}U(x)$. The consumer's monetary benefit is again given by (10), and the consumer maximizes

$$\begin{aligned} \mathrm{E}[U(B_t(k))|\mathcal{H}_t] &= Q_t(k)U(v - j_t(k)) \\ &+ [1 - Q_t(k)]U(k - j_t(k)). \quad (25) \end{aligned}$$

Differentiation, and the identity (8), yield

$$\begin{aligned} &\frac{\partial}{\partial k}\mathrm{E}[U(B_t(k))|\mathcal{H}_t] = \\ &\quad q_t(k)[U(v - j_t(k)) - U(k - j_t(k))] \\ &\quad - Q_t(k)[1 - Q_t(k)][u(v - j_t(k)) - u(k - j_t(k))]. \quad (26) \end{aligned}$$

The choice of $k = v$ clearly causes the derivative (26) to vanish. To see that $k = v$ is the unique choice maximizing (25), recall that $U(\cdot)$ is strictly increasing and $u(\cdot)$ is decreasing. Thus if $v > k$, the first term in (26) is strictly positive, while the second term is nonpositive, so that their difference is strictly positive. A similar calculation holds for the case $v < k$. Thus,

$$\begin{aligned} \frac{\partial}{\partial k}\mathrm{E}[U(B_t(k))|\mathcal{H}_t] &> 0 \text{ if } k < v \\ &< 0 \text{ if } k > v, \quad (27) \end{aligned}$$

so that $k = v$ is the unique maximizer of (25), showing that risk averse consumers will also self-select the desired contract.

Appendix B: Demand Uncertainty

We now assume that v is a random variable, which is not necessarily independent of p_T. Assume that p_T and v have

[1] Reference [4] mentions the concept of "threshold spot price," which is equivalent to our strike price k, to illustrate the allocation achieved under priority pricing, but the menu of priority service contracts is specified as contract price as a function of service probability. The concept of "priority insurance" is much closer to the callable forwards described here [4, 8].

a joint density given by $q_t(p,v)$. We again ask the question of which contract the consumer will select.

As in section 4, we can write the consumer's benefit as

$$B_t(k) = k - j_t(k) + (v-k)1\{p_T < k\}, \tag{28}$$

where $1\{p_T < k\}$ is 1 on the event $\{p_T < k\}$ and 0 otherwise. A risk-neutral consumer maximizes

$$E[B_t(k)|\mathcal{H}_t] = k - j_t(k) - kQ_t(k) + E[v1\{p_T < k\}|\mathcal{H}_t]. \tag{29}$$

Setting the derivative of (29) to 0, substituting (8) and solving for k yields

$$k = \frac{1}{q_t(k)}\frac{\partial}{\partial k}E[v1\{p_T < k\}|\mathcal{H}_t], \tag{30}$$

assuming again that $q_t(k) > 0$, $\forall k > 0$. To evaluate the last term, note that

$$\begin{aligned}
&\frac{1}{q_t(k)}\frac{\partial}{\partial k}E[v1\{p_T < k\}|\mathcal{H}_t] \\
&= \frac{1}{q_t(k)}\frac{\partial}{\partial k}\int_0^\infty\int_0^\infty v1\{p<k\}q_t(p,v)dvdp \\
&= \frac{1}{q_t(k)}\frac{\partial}{\partial k}\int_0^k\int_0^\infty vq_t(p,v)dvdp \\
&= \frac{1}{q_t(k)}\int_0^\infty vq_t(k,v)dv = \int_0^\infty v\frac{q_t(k,v)}{q_t(k)}dv \\
&= E[v|\mathcal{H}_t, p_T = k].
\end{aligned} \tag{31}$$

Thus, the necessary condition for an interior maximum of (29) is

$$k = E[v|\mathcal{H}_t, p_T = k]. \tag{32}$$

When the consumer's valuation is correlated with the spot price p_T, the consumer may "hedge" when selecting a contract. As an example of this, suppose that the consumer's valuation is given by

$$v = v_0 + \beta p_T + \epsilon, \tag{33}$$

where v_0 is a constant, and ϵ is a zero-mean random variable independent of p_T. Since the RHS of (32) is conditioned on the fact that $p_T = k$, the self-selection condition becomes

$$k = v_0 + \beta k, \tag{34}$$

or, if $\beta < 1$,

$$k = \frac{v_0}{1-\beta}. \tag{35}$$

Thus, when $\beta \in (0,1)$, so that the consumer's valuation is positively correlated with p_T, we see that the consumer hedges by purchasing a contract with a higher strike price (higher probability of service) than would be purchased in the case of v and p_T independent. Similarly, if $\beta < 0$, so that there is a negative correlation, then the consumer would hedge lower, purchasing a contract with a lower k than would otherwise be purchased.

Acknowledgements

The authors would like to thank Felix F. Wu, Shmuel S. Oren and Ross Baldick of the University of California, Berkeley, and R. John Kaye of the University of New South Wales, Australia for their helpful comments. This work was funded in part by the National Science Foundation under grants ECS-8715132 and IRI-8902813.

References

[1] Edward Kahn. *Electric Utility Planning and Regulation.* American Council for an Energy-Efficient Economy, Washington, D. C. and Berkeley, CA, 1988. In cooperation with Universitywide Energy Research Group, Berkeley, CA.

[2] M. C. Caramanis, R. E. Bohn, and F. C. Schweppe. Optimal spot pricing: Practice and theory. *IEEE Transactions of Power Apparatus and Systems*, PAS-101:3234–3245, September 1982.

[3] R. J. Kaye, H. R. Outhred, and C. H. Bannister. Forward contracts for the operation of an electricity industry under spot pricing. *IEEE Transactions on Power Systems*, 5(1):46–52, February 1990.

[4] Hung po Chao and Robert Wilson. Priority service: Pricing, investment, and market organization. *American Economic Review*, 77(5):899–916, December 1987.

[5] William F. Sharpe. *Investments.* Prentice-Hall, third edition, 1985.

[6] A. G. Malliaris and W. A. Brock. *Stochastic Methods in Economics and Finance.* Elsevier (North-Holland), 1982.

[7] Chin-Woo Tan and Pravin P. Varaiya. Interruptible electric power service contracts. Interdisciplinary Group on Coordination Theory Memorandum UCB/ERL/IGCT M90/108, University of California, Berkeley, November 1990.

[8] Shmuel S. Oren and Joseph A. Doucet. Interruption insurance for generation and distribution of electric power. *Journal of Regulatory Economics*, 2:5–19, 1990.

[9] Hal R. Varian. *Microeconomic Analysis.* W. W. Norton, second edition, 1984.

[10] Robert C. Merton. Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4(1):141–183, Spring 1973.

Thomas W. Gedra received his B.E.S. and M.S.E. degrees from The Johns Hopkins University in Baltimore, Maryland. He then worked for several years at Westinghouse Electric Corporation, primarily on algorithms for synthetic aperture radar, before attending U.C. Berkeley, where he expects to receive a Ph.D. in Electrical Engineering in May, 1991. His current research interests include stochastic systems and control and power system economics.

Pravin Varaiya is a professor of Electrical Engineering & Computer Sciences at the University of California, Berkeley. He received his Ph.D. in Electrical Engineering from U.C. Berkeley. He is the author, with P.R. Kumar, of *Stochastic Systems: Estimation, Identification, and Adaptive Control*, Prentice-Hall, 1986 and editor, with A. Kurzhanski, of *Discrete Event Systems: Models and Applications*, Lecture Notes in Information Sciences, vol. 103, Springer, 1988. His areas of research and teaching are in stochastic systems, communication networks, power systems and urban economics.