Optimal Expansion of Energy Efficiency Programs: Examples

Allen F. Wilson and Villamor Gamponia, Seattle City Light
Benjamin F. Hobbs, Case Western Reserve University

How quickly should discretionary energy efficiency programs be ramped up? Although such programs have short lead times, significant inefficiencies can occur if they are expanded too quickly. This is because the capacity to deliver DSM cannot be changed instantaneously. Staff must be hired and trained, while non-utility contractors must increase their capability. Adjustment costs also occur later as the resource (e.g., weatherizable homes) is used up. Downsizing too quickly results in hardship for contractors and high employee severance costs.

Optimal program expansion might spread the effort over more years in order to avoid the inefficiencies of a quick run-up and run-down. We present an optimal control framework for scheduling program expansion to maximize net benefits. Among the costs considered is the expense of altering program capacity. The variables that describe program status include staffing level and the number of potential participants that have already participated. The decisions that the utility makes are the number of participants, which is influenced by incentive levels and promotional efforts, and the change in capacity. This control problem is solved using dynamic programming, which can obtain solutions in seconds on a personal computer. A simpler version of the problem can also be solved with a spreadsheet.

As examples, we present applications to the design of weatherization programs for Seattle City Light. Because programs are best spread over several years, we find that it is optimal to initiate programs well before they pass the California Iterative Cost-Effectiveness Methodology "first-year" test.

Introduction

After several years of supply surplus, electric utilities in the U.S. Pacific Northwest suddenly face the possibility of energy shortfalls during the coming decade. In response, the Bonneville Power Administration (BPA) plans to expand its funding of energy efficiency programs several-fold over the next five years. Examples of such programs include commercial lighting audits, free installation of water heater wraps, and subsidized loans for home weatherization. BPA will spend $2.8 billion on that effort over the next ten years. A total of 665 average megawatts of load reduction is anticipated at a cost of $27/MWh in 1990 levelized dollars (BPA 1992).

Seattle City Light (SCL) is a municipally-owned utility that will benefit from the proposed increase in BPA funding. SCL is as well positioned as any Northwest utility to manage growth in its efficiency programs, as it has already invested approximately $100 million in energy conservation. For instance, two-thirds of single family Seattle homes that use electric heat have been weatherized and a major commercial/industrial efficiency initiative is underway.

Yet the increase envisioned by BPA still has the potential to result in administrative inefficiencies at SCL. This is true even if most of the labor is provided by outside energy service companies (ESCOs). SCL's concern over this possible problem motivated this study.

The purpose of this paper is to address the question of how quickly discretionary energy efficiency programs should be ramped up in order to maximize their net benefits.[1] Although such programs have shorter lead times and are therefore often more flexible than new generation sources (Hirst 1989), significant inefficiencies can occur if they are expanded too quickly.[2] The capacity to deliver DSM cannot be expanded or contracted instantaneously. Utility staff must be hired and trained, while non-utility contractors must increase their capability. Too fast an acceleration could lead to unnecessarily high administrative costs early on. In order to induce high participation rates, incentive levels might have to be set at very generous levels, inflating utility costs. Adjustment problems could also occur later as the conservation resource (e.g., weatherizable homes) is used up, resulting in a need to ramp down the program quickly. Too rapid a downsizing results in high severance costs for utility employees and hardship for non-utility contractors who, in turn, are likely to lobby the utility and government for relief.

An optimal expansion of a conservation program might spread the effort over more years in order to avoid the inefficiencies of a quick run-up and run-down. However, no studies have been published which address the question of how quickly programs should be expanded. The Northwest Power Planning Council (1989, p. 1) put it this way:

> *How fast can discretionary conservation programs be sped up ...? How fast will the car accelerate ... if we push the gas pedal to the floor? This issue is plagued with little data. The region has not yet test driven the car at full speed, and thus most of assumptions are judgments.*

The Council has made assumptions concerning the maximum rates at which a program can be ramped up and down, as a percentage of the total potential participants in a program.[3] Their numbers, based on actual experience in the northwest, are rough approximations to the *maximum possible* values--not necessarily the *best* values. Optimal ramp up and ramp down rates are likely to depend on program-specific factors, such as anticipated savings, avoided costs, participation rates as a function of incentives, and the ease with which personnel can be hired and laid off.

In this paper, we present an optimization framework for determining the schedule of program expansion that maximizes net benefits. We solve the most general version of the problem using the technique of dynamic programming; however, solutions for a simplified version can also be obtained using a spreadsheet.

The next section describes our model of the benefits and costs of energy efficiency programs. The subsequent two sections summarize the spreadsheet and dynamic programming methods for solving the problem. Each of the methods is used to optimize a residential weatherization program. Sensitivity analyses show how the optimal expansion of the program depends upon the assumptions made.

The Net Benefits Model

In the version of the model presented here, we assume that the utility's objective is to maximize the net benefit of the program, defined as avoided supply and external costs, minus program costs borne by the utility and program participants. This is equivalent to the widely-used "total resource cost" (TRC) test (CPUC and CEC 1987). It is also possible to maximize a net benefit expression that considers the benefits of additional energy services consumed by participants ("rebound"), benefits of rate changes, and benefits of reliability improvements (Hobbs 1991; Hobbs and Wilson 1992).

Definition of Costs and Notation

We assume that avoided supply costs equal the cumulative number of participants in the program (net of free riders and any attrition) times the energy savings per participant, multiplied by the marginal cost of energy supply. The utility's program costs are modeled in some detail because of their direct impact on optimal ramp up and down rates. The latter costs include the following items:

1. Wages of skilled labor used in program. This labor may consist of utility employees or contractors.

2. Labor adjustment costs, equaling the expense of changing the number of skilled personnel. For instance, hiring expenses might include staff time involved in interviews or contract negotiations, along with wages expended during the "break-in" period for new employees. Staff reduction costs include severance payments and administrative expenses which, for SCL, are likely to be equivalent to one year's salary or more. Even if staff are merely transferred to other programs, there would still be significant retraining expenses.

   These costs are incurred whether the labor in question is hired directly by SCL or through a contract with an ESCO, although their magnitude might vary.

3. Capital cost of installing the energy efficiency measure, including equipment, materials, and participant time. We assume that this is proportional to the number of participants in a given year. Of this cost, only a fraction is borne by the utility, as determined by the incentive it offers participants. However, the total cost paid by both the utility and participants is included in our resource cost objective. Incentives are transfer payments and are not counted in the net benefits expression. (However, if benefits of rate changes are considered, then incentives do affect economic efficiency; see Hobbs 1991.)

4. Other program costs, assumed to be independent of the number of participants. These include marketing costs, fixed management costs, and other overhead expenses.

Below we define the decision variables, functions, and parameters that describe the costs in our model. Decision variables, whose levels the utility can control, are designated by upper case letters. Decision variables are of two types: those which are directly controlled by the utility, and those which are functions of the control variables. Cost functions of those variables are given by lower case letters, as are most fixed parameters, such as interest rates.

Direct Control Variables:

H_t = Net skilled labor hired at the beginning of year t (persons), including contracts with ESCOs. If $H_t<0$, then workers are being laid off or contracts terminated. Labor is considered to be acquired at the start of the period during which it is first productive. The actual hiring must usually take place some time before then, because of the time needed to be trained and become fully productive. Wages and other costs associated with this training period are counted in the cost of hiring.

I_t = The level of incentive offered by the utility to participants, expressed as a fraction of the capital cost cc_t. The incentive can consist of direct monetary payments, decreases in electric bills, or subsidies of materials or labor.

Other Decision Variables (Affected by Control Variables):

L_t = The amount of productive skilled labor available to the program during year t (persons). This can be controlled by altering H_t, as shown below:
= $L_{t-1}(1-a) + H_t$
where a is the work force attrition rate.

N_t = Number of installations of energy efficiency equipment in year t. These installations are treated by the model as if they occur at the beginning of the period.
= $N_t(L_t, I_t, S_{t-1})$. The number of participants in year t is the minimum of the following two quantities: (1) the demand by potential participants for the program, which depends on incentives, the number of potential participants, and the anticipated energy savings; and (2) the capacity of the program, which we measure by the amount of skilled labor. Those two quantities are functions of the decision variables L_t, I_t, and S_{t-1}. If the potential participants exceeds the size of the program, then some participants are forced to wait.

S_t = Cumulative number of participants from time 1 up to and including t, which is controlled by N_t thus:
= $S_{t-1}(1-d) + N_t$
= $\Sigma_{\tau=1,2,..,t}(1-d)^{t-\tau}N_\tau$
where d is the demolition or replacement rate.

Functions:

$c_t(H_t)$ = Total cost of adjustment H_t in skilled labor capacity in year t (\$). This is the cost of hiring workers, if $H_t>0$, and the cost of laying them off should $H_t<0$.

$f_t(N_t)$ = Fixed administrative and marketing costs in year t (\$), incurred only if the program is in operation. It equals 0 if $N_t=0$, while equaling fixed amount f_t if $N_t>0$.

Parameters:

a = Fraction per year of work force that voluntarily quits.

cc_t = Investment cost of the energy efficiency measure, excluding the utility's labor expense (\$/participant). This equals the investment per kWh saved annually times the annual kWh savings. We assume this is constant, although higher incentive levels I_t might encourage participation of consumers for whom energy efficiency would be more expensive.

d = Rate of demolition or replacement of past installations of energy efficiency measures (1/yr). An example is the demolition rate for weatherized homes.

e_t = Energy savings (kWh/participant/yr).

i = Interest rate (1/yr)

mc_t = Marginal cost of energy supply to the utility in year t (\$/kWh)

me_t = Marginal external cost of energy supply in year t (\$/kWh); for Seattle City Light, this results primarily from CO_2 emissions from new supply resources

$s_{max,t}$ = Total population of potential program participants in year t

t = Index of year, or other time period used in model

T = Last year in which the program can be operated

w_t = Wage rate for skilled labor in year t (\$/person-yr)

Net Benefits Model

Using the above notation, the present worth of the program's net benefit becomes:

$$\textit{Net Benefit} = \Sigma_t \{[mc_t + me_t]e_t S_t - \quad \text{[Avoided Supply Costs]}$$

$$[wL_t + c_t(H_t) + f_t(N_t) + cc_t N_t]/(1+i)^t \quad \text{[Program Costs]} \tag{1}$$

This is the discounted sum of avoided supply costs (including externalities), minus utility and participant program costs (including labor, fixed costs, and equipment).

The problem we address in this paper can therefore be phrased as follows: Choose the amount of labor to hire/fire H_t and the level of the participant incentive I_t in each year in order to maximize Equation (1). Large values of H_t and I_t in early years result in a "crash" program; smaller values instead indicate that an energy efficiency measure is to be phased in over a longer period of time. If the functions and parameters have been estimated, then the optimal (i.e., net benefit maximizing) values of H_t and I_t can be solved for by the methods in the next two sections.

An important omission in Equation (1) is the effect of free riders, who can be significant for some programs. Basically, free riders inflate utility expenditures without providing additional savings. In Equation (1), then, energy savings e_t and captial cost cc_t per participant should be multiplied by (1-fr), where fr is the fraction of participants who would have weatherized anyway.

If desired, Equation (1) can be modified to reflect just costs to certain subsegments of society rather than society as a whole. For example, in the second application in this paper, we contrast the effect of using a revenue requirements objective instead of (1).

Simplified Model and Results

We propose two methods for solving the ramp-up problem, Equation (1). For the most general case, in which the functions such as $c_t(H_t)$ may be nonlinear and the utility can control participation rates by altering the incentive I_t, a method called dynamic programming is needed. However, with some simplifying assumptions, the problem reduces to the determination of just two variables: when the program should start and how intensive it should be. A spreadsheet model can then be used to identify good solutions and to gain insight. The simple model is the subject of this section, while in the next section, the full dynamic program is presented.

Optimization of Program Timing by a Spreadsheet

Optimizing Equation 1 becomes easier if the following additional assumptions are made:

1. Labor adjustment costs $c_t(H_t)$ are proportional to the absolute value of the change in labor force H_t ($c_t(H_t) = c_t|H_t|$). There is no attrition of workers (a=0).

2. Incentive payments are adjusted, if necessary, so that participation rates N_t are constant during the life of the program. The program continues until the entire population of potential participants (or some predetermined fraction) has participated. Under these assumptions, I_t does not have to be explicitly accounted for when solving the model.

3. The program will have been completed by year T. To accomplish this, the intensity and timing of the program is adjusted so that the same total cumulative participation results, no matter what program design is adopted. Further, attrition is disregarded (d=0). These assumptions allow us to ignore benefits and costs after year T, since the savings after that year will be constant, unaffected by the decisions made before T.

The result of these assumptions is that participation in the optimal program generally has the shape shown in Figures 1b and 2b. Participation rates are the same during every year of the program, except the last year. All the labor needed by the program is hired in the first year of the program, and all the labor is laid off during the last or next-to-last year. Consequently, the problem reduces to one of timing (when should the program start?) and intensity (how many participants should there be per year?). By simply searching over the range of possible values of those variables, and calculating the resulting performance of the program, an optimal configuration can be obtained.

Example Applications of Simple Model

The effect of the timing and intensity decisions in the simple model is shown by Figures 1 and 2. They show the results of two possible configurations of a generic weatherization program based on recent work at SCL on multifamily buildings. Most of the energy savings are due to installation of energy efficient windows, with some additional savings from lighting changes and insulation.

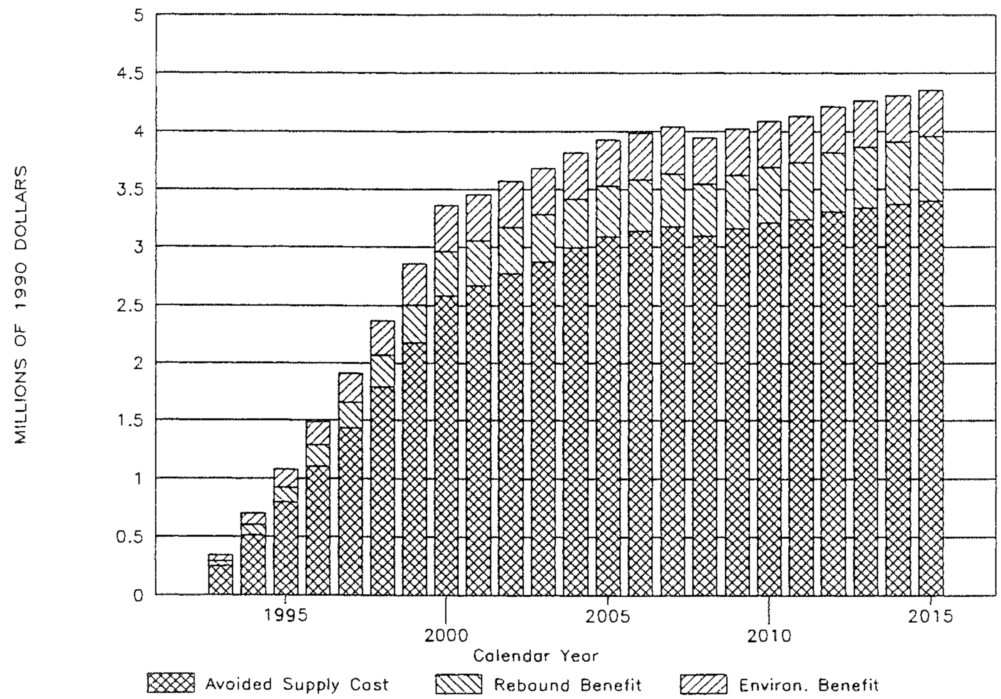

(a) Environmental, rebound, and avoided supply cost benefits

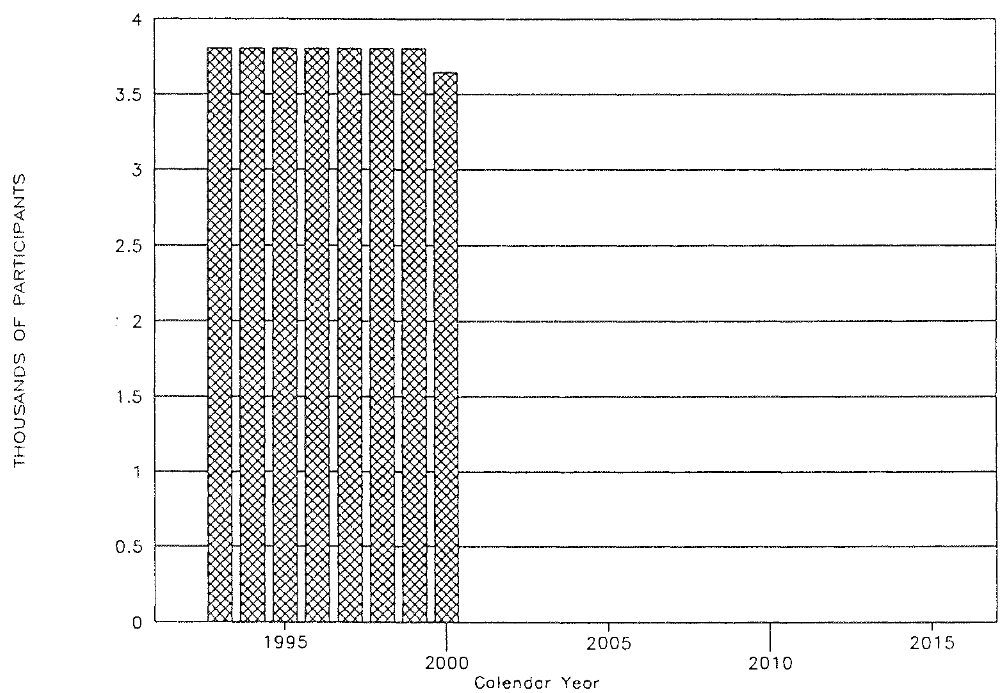

(b) Participation rates, assumed to be proportional to labor

Figure 1. Yearly Program Benefits and Participation for Example Multifamily Energy Efficiency Program

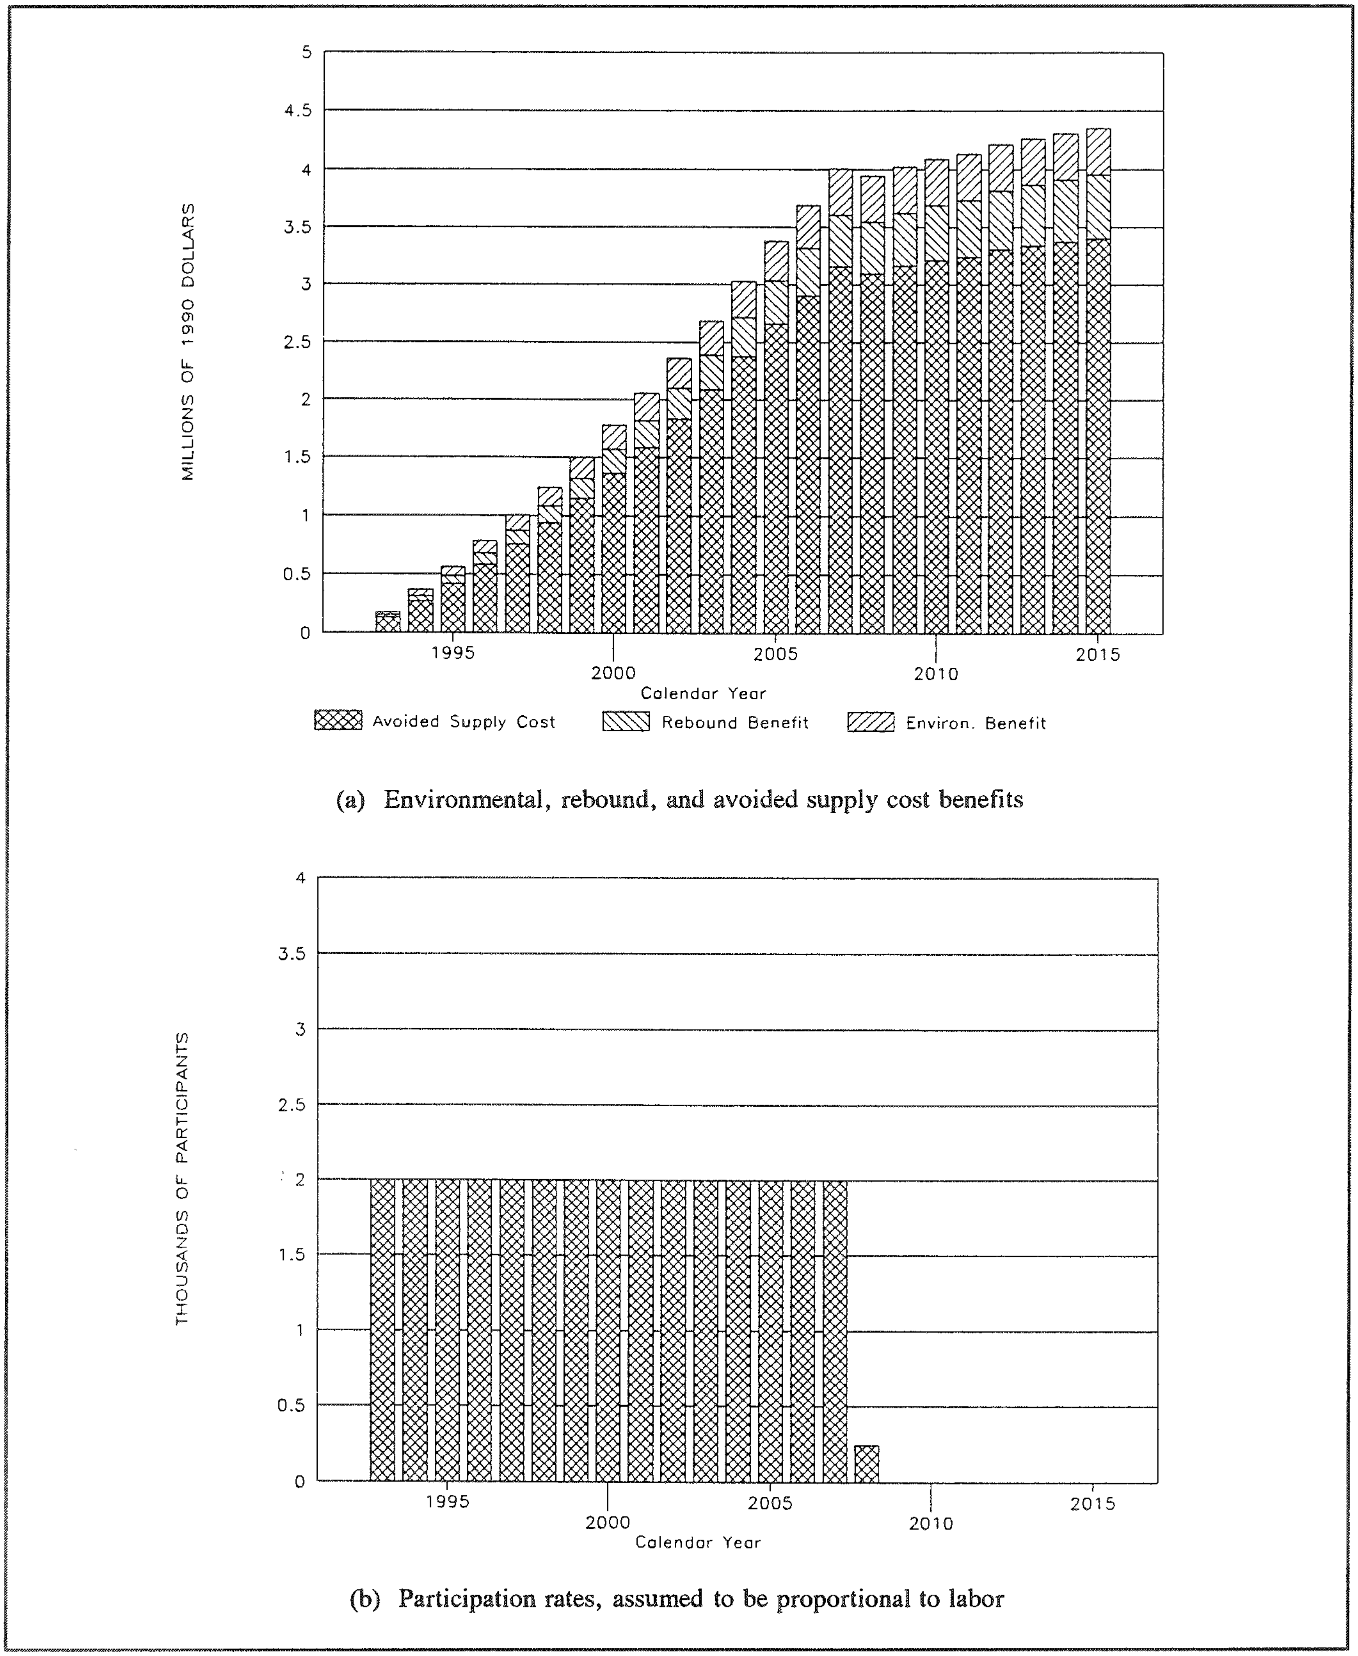

Figure 2. Yearly Program Benefits and Participation for Alternative Multifamily Energy Efficiency Program

Figures 1a and 2a show the benefits resulting from the programs in each year, including avoided supply and environmental costs and benefits of rebound.[4] The first configuration is an intensive program that is started immediately and is completed in eight years. The second configuration is a program which is continued over a longer period of time at a less intense level. The first configuration results in greater labor adjustment costs than the second configuration, because the expense of hiring, training, and firing is greater. However, the first configuration also yields more benefits more quickly, as a comparison of Figures 1a and 2a show. Which configuration is optimal depends on whether the increased adjustment costs are justified by the additional benefits received.

Additional runs show that in the extreme case in which there are no labor adjustment costs, the optimal decision is to have a very intensive program that installs all the measures in one year. The optimal year to implement the program would be the first year in which the benefits of the program in that year exceeded the annualized cost of the program--the so-called "crossover year." This criterion is analogous to the first-year test of the California Iterative Cost Effectiveness Methodology (Kirshner 1992).

However, if adjustment costs are significant, the program should be spread over several years and *it is instead optimal to start the program earlier.* Waiting to start until the year of crossover would result in foregoing of some benefits by some participants, without a concomitant savings in program costs. If the benefits per kWh saved of energy efficiency are increasing linearly over time, it turns out that the best timing is to start the program so that slightly less than half of the participants have had the measure installed by the crossover year.

SCL has used this general approach to optimize the timing and intensity of a multifamily residential weatherization program. The result of the analysis was a recommendation that program effort be trebled immediately.

Strictly speaking, however, the conclusions made above are valid only if the three simplifying assumptions described earlier in this section are adopted. Because those assumptions may be inappropriate, and because we may also be interested in optimizing the incentive level over time, the dynamic program described below is needed.

Dynamic Programming (DP) Model Application

Dynamic programming is a powerful means of solving optimization problems involving the allocation of resources over time (Hillier and Lieberman 1990). Elsewhere, we describe exactly how the method is applied to our problem (Gamponia et al. 1992). The method yields the optimal values of the variables H_t, I_t, L_t, N_t, and S_t in each time period t, and the resulting net benefits (Equation 1). Below, we summarize an application of this dynamic programming methodology to a hypothetical residential weatherization program.

DP Case Study: Assumptions

A pool of 40,000 single family residences using electric heating is the target of the weatherization program examined here. Table 1 summarizes our numerical assumptions. Because our purpose is to demonstrate the methodology, many of the assumptions are simplistic. Some of the most important assumptions are:

- The cost of hiring or laying off one person is \$50,000, equal to one year's wages, benefits, and indirect labor costs. We assume that the utility bears this cost whether the labor involved is hired directly by the utility, or indirectly via an ESCO. If the program is in operation, a fixed management and marketing cost of \$250,000/year is incurred.

- The number of participants of the program is the smaller of the following two quantities: the number of participants that the existing labor L_t can service and the number of homeowners who respond to the incentive I_t. One person is assumed to be able to weatherize forty homes per year. A simple linear relationship is assumed between homeowner response and the incentive. If the incentive in a year t covers 100% of the investment cost cc_t, then we assume that 40% of the remaining unweatherized homes would potentially participate in that year.[5] Lower incentives cause potential participation to drop proportionally. The utility is allowed to vary the incentive rate over time to optimize participation.

- There is no attrition (d=0). Any home that is weatherized is assumed to result in energy savings indefinitely. No additional homes are weatherized after year 20.

- Benefits due to rebound, rate changes, reliability improvements, and avoided environmental costs are disregarded.

- The marginal cost of energy is presently \$0.05/kWh and grows at 2%/year.

Five cases are examined. A base TRC case uses the above assumptions and maximizes the present worth of

Table 1. Dynamic Programming Case Study Assumptions (All Values in Real Dollars)

N_t	=	$\mathrm{MIN}[L_t/40, 0.4I_t(s_{max,t} - S_t)]$ homes weatherized
$c_t(H_t)$	=	$50{,}000\lvert H_t \rvert$ \$
$f_t(N_t)$	=	250,000 \$ if $N_t > 0$
a	=	0/year
cc_t	=	\$1000/home weatherized
d	=	0.0/year
e_t	=	3000 kWh/yr/home weatherized
i	=	0.03/yr
mc_t	=	$0.05e^{0.02t}$ \$/kWh
me_t	=	0.0 \$/kWh
$s_{max,t}$	=	40,000 homes
T	=	20 years
w_t	=	50,000 \$/person-yr

avoided supply costs minus program costs (essentially, the TRC criterion). The second case, the **naive TRC case**, ignores labor adjustment costs in determining the optimal timing and size of the program. This case allows us to show the potential resource loss that results from disregarding the expense of ramping up. Then, a **base revenue requirements case** is run in which the utility minimizes the present worth of revenue requirements instead of maximizing net benefits. The difference between that and the base TRC case is that the capital cost per person is $I_t cc_t$ in the revenue requirements objective, rather than the full cost cc_t. The fourth case is a **naive revenue requirements case** which, like the naive TRC case, is developed ignoring adjustment costs. The fifth and final case is the **fixed incentive revenue requirements case** in which the incentive is assumed to be fixed at $I_t = 0.5$ during the program, rather than being allowed to vary. This means that the only control variable available to the utility is the number of workers hired or laid off.

DP Results

Table 2 gives the results for each of the cases. Two indices of program net benefits are shown: (1) the present worth of societal net benefits (Equation 1) and (2) the present worth of the decrease in revenue requirements. The former is always the smaller of the two because it includes all installation costs, not just those borne by the utility. We also show when the program is started, how long it lasts, the range of incentives that are offered, and the largest number of skilled workers employed by the program.

Our major conclusion is that ignoring adjustment costs in program design can result in a significant loss of benefits. Generally, this is because the naive cases are more intensive and result in greater numbers of workers being hired and then almost immediately laid off. For instance, in the base TRC case, 150 workers are hired in year 0, and then 50 workers are laid off in years 4, 5, and 6. However, the naive TRC case hires 400 workers right away, lets half of them go the next year, and lays the remaining workers off over the remaining life of the program. As a result, the net societal benefits fall by one-third (\$20 million = \$60.9 - \$40.8 million) compared to the base TRC case which optimized hiring/firing.

An additional benefit of the less intensive program in the base TRC case is that lower incentives can be offered, thus diminishing impacts on nonparticipants. This is reflected in the difference of \$30 million in revenue requirements between the base and naive TRC cases (\$72.9 - \$42.8 million), which is 50% greater than the difference in social net benefits.

If revenue requirements are to be minimized, rather than social net benefits maximized, the utility is motivated to lower the incentives and spread out the implementation of the program. The result is lower social net benefits, as the foregone energy savings are more valuable than the economies resulting from using a smaller work force. However, the loss of social net benefits is only about \$1 million (\$60.9 - \$59.9 million) compared to the Base TRC case, a negligible amount.

If adjustment costs are ignored, then the revenue requirements objective leads to a smaller loss of social net benefits than the TRC objective. The difference between the net benefits of the base and naive revenue requirements cases is only \$5 million (\$59.9 - \$54.9 million), a fourth of the loss that occurs if instead the utility tries to maximize social net benefits (\$60.9 - \$40.8 million). The reason is that even if the utility disregards adjustment costs, it is still motivated to keep the program small so that incentives (and thus revenue requirements) won't have to be as high.

In Cases 1 through 4, the utility has the ability to vary the incentive I_t in order to control participation. As a result, optimal incentive rates are generally low and then rise with time as the population of unweatherized homes shrinks. Then, as the program lets workers go in its waning years, the incentive might be decreased. Case 5 examines the effect of removing that control by assuming that the incentive is fixed at 0.5 of the measure's capital cost. The consequence is more than a 10% fall in social net benefits (from \$60.9 million to \$53.4 million), in large part because fewer people are enticed to participate.

Table 2. Results for Hypothetical Residential Weatherization Program

Case	Present Worth of Societal Net Benefits ($1,000,000)[a]	Present Worth of Decrease in Revenue Requirements ($1,000,000)[a]	First, Last Year of Program	Smallest, Largest Value of I	Maximum Number of Workers
1. Base TRC	60.9	72.9	0,6	0.38,1.0	150
2. Naive TRC	40.8	42.8	0,4	0.83,1.0	400
3. Base Revenue Requirements	59.9	78.3	0,10	0.25,0.83	100
4. Naive Revenue Requirements	54.9	70.3	0,8	0.42,0.83	200
5. Fixed Subsidy, Revenue Requirements	53.4	67.8	0,9	0.5,0.5	100

(a) Includes adjustment costs in all cases, even though the naive cases have ignored those costs when designing the program.

Cases 1 through 4 weatherize 90% of the eligible homes, but Case 5 weatherizes only 80% because it cannot raise the incentive in later years. Thus, the ability to manipulate the incentive level can bestow benefits upon both the utility and society.[6]

Finally, we also made runs with other values for energy savings and the worth of those savings. We found that for marginal programs, it may be optimal to delay the program until later when its worth has risen high enough to justify it. If adjustment costs are ignored, it is optimal to maximize participation at that time. But in the presence of the constraint, it is actually better to start the program earlier in order to minimize adjustment costs and to capture the benefits of the savings more quickly, as we pointed out earlier in this paper.

Conclusion

In general, it is costly to adjust the capacity of a DSM program in large part because of the expense of hiring and laying off labor.[7] Consequently, an optimal program is spread over a greater number of years than it would be otherwise. The program is also started earlier in order to capture benefits of energy savings that would otherwise be foregone. These conclusions result from both the spreadsheet and dynamic programming analyses.

Further work is needed on the following issues:

- What are the implications for timing of multiple DSM programs? If timed carefully, labor might be transferred among the programs at less cost than the expense of hiring new workers.
- What is the effect of uncertainty in future avoided costs and participation rates? In particular, what if the program size could be adjusted over time in reaction to what is learned about demand growth and program effectiveness (Hirst 1991)?
- What is the effect of a diverse pool of potential participants among which there are a variety of costs, potential savings, and predilections to participate?

Endnotes

1. We assume that installation of energy efficiency measures can be deferred at no cost, except for the foregone energy savings. This implies that programs that address "lost opportunity"-type resources, such as new commercial buildings, cannot be analyzed in this framework, as delay means that the opportunity for installing inexpensive DSM measures at the time of construction is lost.

2. For instance, the Northwest Power Planning Council (1989) notes that in 1983, when more than 89,000 houses were weatherized in the Northwest, "the speed with which things were happening led to some disarray" (p. 2). So the Council suggests that 50,000 households per year, the number weatherized in 1984 through 1986, "might be a reasonable maximum" (*ibid.*).

3. For example, existing residential space heat programs can reach a maximum of 14% of the potential participants per year; the program size can be changed up or down in a year by no more than 7% of the potential

participants. These values are much smaller for other programs.

4. Rebound benefits equal the increase in participant consumer surplus resulting from changes in temperature settings (see Hobbs 1991 or Wilson and Gamponia 1990). The figure points out the importance of considering changes in customer value, as the estimated rebound benefits are similar in size to the avoided environmental costs.

5. This figure is based on experience in the Hood River project when weatherization was provided at no cost to the participant.

6. However, this benefit may be exaggerated. The reason is that in reality the program is less likely to be cost-effective for the later participants because their costs would probably be higher and their potential savings lower than earlier participants.

7. E. Kahn (personal communication) points out that another adjustment cost is the possible loss of utility credibility that could result from fluctuations in program effort and incentives. The consequence may be a diminishing of the willingness of potential participants to participate in the program.

References

Bonneville Power Administration. 1992. *1991 Pacific Northwest Loads and Resources Study*. Portland, Oregon.

California Public Utilities Commission and California Energy Commission. 1987. *Standard Practice for Cost-Benefit Analysis of Conservation and Load Management Programs*, San Francisco, California.

Gamponia, V., B. F. Hobbs, and A. F. Wilson. 1992. "Methods for Optimizing the Expansion of Energy Efficiency Programs." Submitted for publication.

Hillier, F. S., and G. J. Lieberman. 1990. *Introduction to Operations Research*, 5th Ed. McGraw-Hill, New York.

Hirst, E. 1990. "Flexibility Benefits of Demand-Side Programs." *The Energy Journal*, 11(1):157-165.

Hobbs, B. F. 1991. "The 'Most Value' Test: Economic Evaluation of Electricity Demand-Side Management Considering Customer Value." *The Energy Journal*, 12(2):67-91.

Hobbs, B. F., and A. F. Wilson. 1992. "Most Value Planning: Acquiring Supply- and Demand-Side Resources to Maximize Net Benefits." Paper presented at the *American Public Power Association Engineering and Operations Workshop*, March 10, San Antonio, Texas.

Kirshner, D. 1992. "Practical Methods for Capacity Expansion: Lessons Learned (?) in California." Paper presented at conference on *Electric Markets and All Resource Options: Beyond Integrated Resource Planning*, San Francisco, California, Feb. 27.

Northwest Power Planning Council. 1989. *Three Key Conservation Assumptions: Conservation Flexibility, Achievable Conservation and Residential Standard Operating Conditions*. Staff Issue Paper 89-26, Portland, Oregon.

Wilson, A. F., and V. Gamponia. 1990. "Market Demand for Seattle's Weatherization Program: An Economic Assessment." Summarized in *Proceedings: Innovations in Pricing and Planning*. EPRI CU-7013, Electric Power Research Institute, Palo Alto, California.