Where Did the Money Go? The Cost and Performance of the Largest Commercial Sector DSM Programs

*Joseph Eto, Suzie Kito, Leslie Shown, and Richard Sonnenblick**

Spending on electricity energy-efficiency programs was responsible for most of the growth (and decline), and almost all of the energy savings from U.S. utility demand-side management (DSM) programs between 1990 and 1998. As a result of restructuring, utilities may never again assume such an important role in promoting electricity energy efficiency. However, as governments consider future domestic policies to promote energy efficiency in response to global environmental commitments, the potential of large-scale energy efficiency programs will likely be discussed. This article presents new information on a critical issue that will surely arise in these discussions: how much does it cost to save energy through programs that use monetary incentives and targeted information to influence individual customer decisions? We present findings from a detailed examination of the complete costs and measured energy savings from the largest commercial sector DSM programs operated by U.S. electric utilities in 1992. We extend the methodological considerations first identified by Joskow and Marron (1992) regarding differences among utility cost accounting conventions and savings evaluation methods. We quantify the impact of missing and incomplete data and, to the extent they can be assessed, demonstrate that our assumptions to address them are conservative in that they err on the side of overstating the apparent cost of saved energy. We find that the programs, as a whole, have saved energy at a cost of 3.2¢/kWh. When compared to the cost of the energy they allowed the sponsoring utilities to avoid generating or purchasing (in the absence of these programs), we find that the programs, as a whole, are cost effective.

INTRODUCTION

Between 1990 and 1998, U.S. electric utilities spent over $18 billion, reaching and annual peak of nearly $3 billion in 1993, on demand-side management (DSM) programs to actively influence their customers' use of energy (EIA, 1999). By 1998, annual spending on DSM had fallen by about half

The Energy Journal, Vol. 21, No. 2.

* Lawrence Berkeley National Laboratory, MS 90-4000, 1 Cyclotron Road, Berkeley, CA 94720. E-mail: jheto@lbl.gov

to $1.6 billion with spending projected to remain constant as the industry undergoes structural changes to accommodate increased wholesale and retail competition. Spending on energy-efficiency programs was responsible for most of the growth (and decline), and almost all of the energy savings from utility DSM programs during this period. The programs represent an unprecedented level of spending on activities to promote the adoption of energy-efficient technologies and practices. Most agree it is unlikely that U.S. electric utilities will ever again assume such an important role in promoting electricity energy efficiency.[1] However, as governments consider future domestic policies to promote energy efficiency, for example, in response to global environmental commitments, we expect that the potential role for large-scale energy efficiency programs will be discussed.

Key questions for discussions of future, large-scale energy-efficiency programs include: how much energy can be saved through programs that influence individual customer transactions with monetary incentives and targeted information and, more importantly, how much do they cost? Sadly, current information on utility-sponsored energy-efficiency programs is incomplete.

Joskow and Marron (1992) broke important new ground by systematically describing a framework for organizing the issues associated with measuring the cost of energy saved by utility programs. They demonstrated the importance of accounting for all costs and identified some of the issues associated with measuring energy savings. They also found that utility accounting methods and savings evaluations were generally inconsistent with one another, making it difficult to assess the total cost of energy saved by the programs. For example, the U.S. Energy Information Agency annually collects information on DSM program costs and projected energy savings from electric utilities (EIA 1997). However, utility cost data, alone, are an incomplete measure of the total social cost of energy savings because they do not include the additional, out-of-pocket costs contributed by the customers that participate in utility DSM programs. Moreover, by the nature of the reporting format, the reported energy savings represent utility forecasts based on current year expenditures, not savings that have been verified following installation.

Parfomak and Lave (1996) provide a partial answer to the question of whether the programs saved energy by demonstrating that the energy savings from utility DSM programs can be identified in aggregate through a macro-level examination of utility sales data. However, their study did not seek to determine

1. While utilities, per se, may be less involved in promoting energy efficiency, many states have adopted and others are considering a surcharge on electricity sales to continue ratepayer funding for energy efficiency, possibly involving non-utility administration. See, Eto, Goldman and Nadel (1998).

the cost of these savings or the relative contribution of different types of DSM programs.

Eto et al. (1996) applied Joskow and Marron's framework to 20 utility commercial lighting DSM programs offered by 18 utilities and addressed many of the shortcomings Joskow and Marron had observed with regard to incomplete reporting of costs and variations in savings evaluation methods. They concluded that the programs had saved energy at a cost of 3.7¢/kWh. While important from a methodological perspective, the findings from the small number of programs analyzed in their convenience sample does not support broader generalizations on the cost of energy savings.

In summary, we are unaware of any study that has comprehensively examined a significant portion of utility spending on DSM to determine the cost of saved energy. This article is a contribution to this deficiency. In it, we present findings from a detailed examination of the costs and measured energy savings from the largest commercial-sector utility DSM programs operated in 1992. Taken together, the programs account for a significant fraction of total U.S. utility spending on energy efficiency in that year. We offer two modest extensions to the methodological framework outlined by Joskow and Marron (1992). Most important of all, we directly assess the effects on the cost of saved energy of: (1) key cost accounting omissions; (2) different savings evaluation methods and assumptions regarding the persistence or longevity of savings.

The article is organized in six sections following this introduction. In section 2, we briefly review the 40 commercial-sector utility DSM programs we examined. In section 3, we describe the extensions we have made to Joskow and Marron's original discussion of DSM program costs in which we clarify treatment of measure costs, and discuss our rationale and methods for including utility administrative costs and utility shareholder incentives. In section 4, we describe the methods we developed to account for missing cost data and assess the direction and magnitude of the biases we may have introduced. In section 5, we review the evaluation methods used by the utilities to measure energy savings and the assumptions utilities made regarding the longevity of savings. As in the previous section, we also discuss the influence of differences in evaluation methods and longevity assumptions on our findings. In section 6, we present our findings on the cost of energy saved by the programs and on the cost-effectiveness of the programs. In section 7, we report findings from our preliminary efforts to explain variations in the costs of saved energy. Section 8 contains our conclusions.

THE LARGEST COMMERCIAL SECTOR DSM PROGRAMS

We focused on programs targeted to the commercial sector programs in this study because the energy-efficiency opportunities there are thought to be large and highly cost effective.[2] As a result, commercial sector programs often represent the largest single element in a utility's portfolio of DSM programs. We focus on 1992 programs because post-program evaluations for 1992 were the most recent ones consistently available when we began our study. We sought information on only the largest commercial sector DSM programs—those with a budget of $1 million or more—so that our results would capture a substantial fraction of utility DSM spending in 1992.[3] Several of the programs were among the largest DSM programs in 1992.

Utility spending on the 40 programs in this study ($380 million) represents nearly a third of total 1992 industry spending on energy-efficiency DSM programs ($1.2 billion). The programs accounted for more than half of the total energy efficiency DSM program budget of the 23 sponsoring utilities ($720 million). The total spending by these sponsoring utilities, in turn, accounted for over 2/3 of total industry spending on energy-efficiency DSM programs in 1992; not only are the programs we examine among the largest in the industry, the sponsoring utilities, themselves, are also industry leaders in total DSM spending.

The programs were all full-scale programs (as opposed to pilots), but varied in maturity. Five only began full-scale operation in 1992, while three began full-scale operation prior to 1986. The majority of programs (29) were rebate programs, while the remaining programs were direct installation programs (11). We did not include programs that targeted only new construction, although some of the programs we examined also offered rebates for equipment upgrades in buildings under construction (or, more typically, undergoing renovation).

We categorized the majority of programs as multi-measure programs; the next largest number were lighting-only programs (see Figure 1). While multi-measure programs promoted measures for all major commercial sector end uses, including lighting, heating, ventilating and air conditioning (HVAC), motors, building shell or envelope, refrigeration, water heating, and process, lighting measures from these and the lighting-only programs accounted for the majority of the savings from all of the programs. The lighting technologies promoted by the programs were quite similar. For 30 of the 35 lighting and multi-measure programs for which we had information, 26 promoted compact fluorescent lamps, electronic ballasts, and either T-8 or T-12 fluorescent lamps; 24 promoted reflector systems; and 22 promoted lighting controls and high intensity discharge lamps.

2. Looking forward, recent studies continue to suggest that substantial cost-effective energy savings opportunities remain in this sector (IGWELT, 1997).

3. See Eto et al. (1995) for a detailed description of the data collection process.

Figure 1. DSM Program Types and Annual Savings by End Use

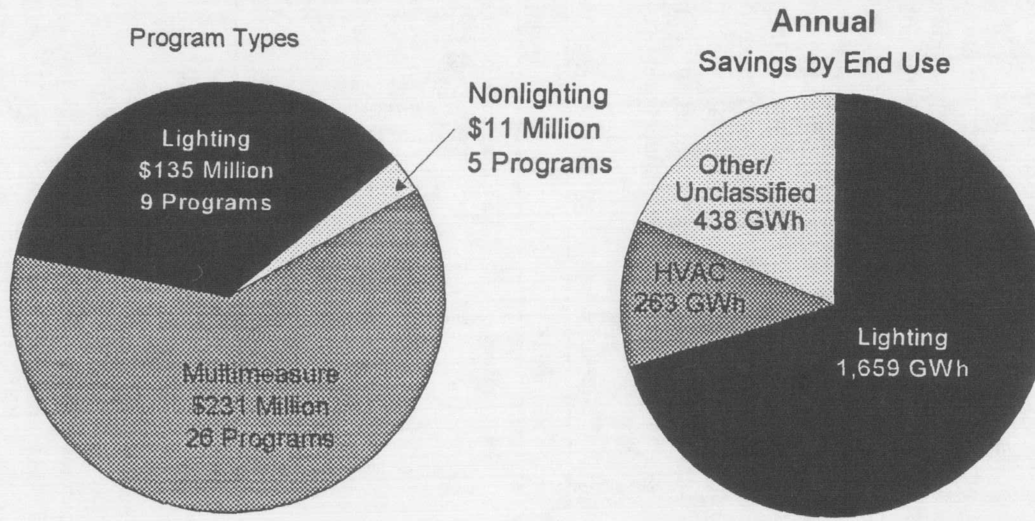

Note: Total utility spending on the 40 programs is $377 million or about one-third of total industry spending on energy efficiency DSM programs in 1992. Total annual savings are 2,360 GWh.

EXTENDING JOSKOW AND MARRON'S ANALYTIC FRAMEWORK

Joskow and Marron (1992) outline a comprehensive framework for measuring the total cost to society of energy saved by utility DSM programs, which we refer to as the total resource cost (TRC) of saved energy. The total resource cost of saved energy consists of two types of costs, measure and non-measure costs (Table 1). Measure costs are the total cost of purchasing and installing the energy-efficiency measures promoted by the DSM program. Measure costs may be borne by the participating customer, the utility (e.g., a rebate), or, more typically, by both (e.g., a rebate pays only a portion of the total cost of a measure). Non-measure costs are borne only by the utility. They are the costs associated with operating a program to promote the adoption of energy efficiency measures (excluding the cost of the measures, themselves), such as advertising expenses, evaluation costs, and administrative overheads. Including both types of costs is essential for assessing accurately the total cost of saved energy to society. In this section, we describe two modest enhancements that 1) clarify important definitional issues associated with the treatment of measure costs and 2) discuss the inclusion of a previously unaccounted for element of non-measure costs, shareholder incentives.

Table 1. DSM Program Total Resource Cost Elements

Cost Type	Cost Elements	Who Incurs Them
Measure	Net installation	Program participants and/or Utility
	Net equipment	
Non-measure	Program administration (incl. overhead)	Utility only
	Program evaluation	
	Shareholder incentives	

Identifying the Baseline in Assessing Net Measure Costs

Developing measures costs consistently is complicated by two related issues. First, as noted, measure costs may be borne by the utility, the participating customer or, more typically, partially by both. The measure costs borne by the customers participating in a DSM program are a major element that Joskow and Marron noted was often missing from discussions of the cost of energy savings. Second, depending upon the baseline situation/condition assumed by the utility, only a fraction of the total installed cost of a measure (paid either by the utility or the customer) may be attributable to the energy savings from that measure. For clarity, we refer to these costs as net measure costs. Our analysis sheds light on the importance of identifying the baseline used in assessing net measure costs.

On the one hand, if the baseline situation is that equipment is at the end of its economic life, the decision to replace it is assumed to be imminent with or without the DSM program. For energy-saving measures installed in these circumstances, the incremental measure cost is the difference between the cost of the equipment that would normally replace the equipment being retired and the actual cost of the equipment promoted by the utility. In this situation, net measure costs legitimately may be quite small; for example, there may be no additional installation costs, and additional equipment cost may be only a fraction of the total equipment cost. We call adoption of energy-saving measures at this point in the equipment life cycle "normal" replacements. Normal replacement is common for HVAC measures, in which equipment at the end of its useful life is replaced by new, energy-efficient equipment.

On the other hand, for customers whose equipment is not at the end of its economic life when the DSM program supports its replacement, the net measure cost of energy savings is the full cost of the DSM measure, including the cost of the equipment and its total installation costs, less the salvage value of the equipment, if any.[4] If the customer had not decided to participate in the utility program, no costs would have been incurred (and no savings would have been generated). We call the adoption of energy savings measures in these situations "early" replacements. Early replacement is more common for lighting measures, when working lighting equipment is removed and replaced with energy-efficient equipment, than it is for HVAC equipment.

Although the distinction between normal replacement and early replacement is easy to make in theory, it is often difficult to apply in practice. Nonetheless, it has great implications for the TRC. For both normal replacements and early replacements, savings may or may not be affected. However, because the net measure costs attributable to the savings differ, the TRC will differ even though there is no difference in the total out-of-pocket cost of measures installed. As discussed later in this article, utilities do not consistently report the baseline assumed in reporting measure costs and, consistent with Joskow and Marron's findings, sometimes did not record participant-paid measure costs. We developed a conservative approach to impute

4. In principle, the salvage value of the equipment that is being retired early should be credited against these costs. However, we found no information on the value of salvaged equipment and, hence, did not include such a credit in our calculations. Including such a credit would have lowered the cost of saved energy compared to the findings reported in this paper.

missing data on participant-paid measure costs that assumed all measures were early replacements; this approach has the effect of consistently assuming higher measure costs (and, hence, may introduce an upward bias in the total cost of saved energy).

Inclusion of Shareholder Incentives as a Cost to Society

There are differences of opinion about whether shareholder incentives should be included when estimating the TRC. Some argue that shareholder incentives are no more than transfer payments from ratepayers to shareholders and, therefore, are not a cost to society. However, others argue that shareholder incentives are a cost to society, like management fees, and therefore should be included in the TRC. The difficulty in assessing these positions is that there is no standard for an appropriate management fee for utility delivery of energy savings. Eto et al. (1998) argue that one must posit the existence of "hidden utility costs" in order to justify and establish an appropriate management fee. At the same time, they concede that there are substantial practical difficulties in estimating hidden costs with precision. Moreover, they speculate that the range in current shareholder incentive payments likely exceeds the range of hidden cost. Thus, in economic terms, some of these payments are just transfers.

In view of the difficulty of assessing the dividing line between transfers and management fees, we opted to include the total cost of shareholder incentives in calculating the cost of saved energy. This decision, again, may introduce an upward bias in our assessment of the total cost of saved energy.

DEVELOPING CONSISTENT INFORMATION ON THE COST OF ENERGY SAVINGS

One of Joskow and Marron's key findings was absence of uniform cost accounting definitions and practices among utilities. The most significant challenge we faced was development of methods and procedures to account for differences in accounting practices and in some cases cost information not reported by the utilities. While a comprehensive assessment of the bias that may have been introduced by our procedures to address these differences is not possible given these data, this section reviews our efforts and assesses the magnitude and direction of bias introduced for individual data elements. We first discuss measure costs incurred by the utility and participants in the utilities' programs, and then utility non-measure costs.

Measure Costs

Utility-paid measure costs (rebates and direct installation costs) were generally well-documented for the programs in our sample. However, participant-paid measure costs were not reported uniformly. For 30 of our 40 programs, we obtained a direct estimate of participant-paid measure costs, which were typically based on customer invoices from completed installations. For 10 of our 40 programs, we worked from information on rebate design, program planning filings, and rebate levels to develop an estimate of these costs. For example, utilities often design rebate levels to pay for an assumed fraction of measure costs.

We adopted a conservative approach for treating measure costs that has the effect of overstating the true cost of saved energy. We assumed that, unless otherwise indicated, reported measure costs represented net measure costs. That is, we assumed that the utilities had taken the program baseline into account when developing these costs, even though we had evidence to suggest that utilities had simply reported total measure costs (and had not in fact accounted for a program baseline). In addition, when imputing measure costs from supporting documentation, we assume a program baseline of early replacement that led to inclusion the total costs of the measures (rather than assume a program baseline of normal replacement, which would have the effect of lowering these costs).

One way to assess the effect of the bias that we may be introducing is to examine the difference in the measure cost component of cost of energy savings among programs for which the utilities had explicitly indicated a program baseline had been taken into account. To control for differences in the costs of different types of measures (e.g., HVAC versus lighting), we conducted this examination only for those programs in which lighting accounted for more than 90% of savings.

Our examination indicates that the differences between the two program baselines can have a large effect on the measure cost component of the cost of energy savings. Given the small number of programs (3) that reported incremental measure costs, however, the statistical significance of this finding is at best only suggestive. We can, however, confirm that our efforts to infer missing participant costs based on full measure costs may introduce a bias that has the effect of overstating the total resource cost of saved energy.

Non-Measure Costs

Non-measure costs, which are borne exclusively by the utility, were not surprisingly consistently available for our analysis. Non-measure costs closely related to program implementation were readily identifiable, although we opted to suppress some of the underlying sub-categories used by the utilities due to differences in the sub-categories reported by the utilities. Other non-measure costs, including program overhead costs, evaluation costs, and shareholder incentives required special attention.

Overhead was sometimes reported separately and sometimes included as part of direct costs. As with the other direct cost categories, we wanted to assure ourselves that some allocation of overhead was included. Overhead costs were reported separately for nine programs. For another 29 programs, program documentation or utility staff indicated that overhead was already included in reported direct costs. For another two programs, we could not locate an explicit overhead cost category or determine whether overhead was already included in the direct costs reported.

Some insight into the effect of including (or inadvertently excluding) overhead can be gained by examining the subset of programs for which these costs were explicitly reported. For the nine programs that reported explicit overhead costs, overhead costs averaged 4% (standard deviation 4%) of total utility costs (measure + non-measure costs). We conclude that although overhead costs should be included for the sake of completeness, it does not represent a significant fraction of the total resource cost of saved energy.

Developing information on measurement and evaluation costs presented a different challenge. First, measurement and evaluation costs were sometimes not separately reported but were included in other program cost categories. This was especially true for programs whose primary source of estimated savings information was program tracking databases (discussed below). Second, measurement and evaluation costs reported in program year 1992 generally referred to measurement and evaluation activities conducted to estimate savings from a prior program year. Third, when measurement and evaluation costs were separately reported, they were commonly reported as an aggregate total for all measurement and evaluation activities for a given program year.

Our approach to measurement and evaluation costs was as follows. We generally attempted to identify and report measurement and evaluation costs expended to evaluate savings for the 1992 program year by searching records to locate the future year in which they were incurred (we found them for 14 programs). More commonly, we simply relied on 1992 expenditures for measurement and evaluation of previous program years as a reasonable proxy for the measurement and evaluation costs associated with evaluating the 1992 program (we did this for 23 programs). For three programs, measurement and evaluation costs were not reported separately but were included in another cost category (which was already included in our analysis).

For the 37 programs with identified measurement and evaluation costs, measurement and evaluation averaged 3% (standard deviation 2%) of total utility costs (and an even smaller percentage of the total resource cost). We conclude that, although these costs should be included for the sake of completeness, they also do not represent a significant fraction of the total resource cost of saved energy.

As discussed previously, we include shareholder incentives in our estimate of the total resource cost of energy savings. For utilities that receive

DSM shareholder incentive payments, we were generally able to locate these payments in regulatory filings. However, because of the design of the incentives, the filings typically contained a single amount reflecting the utility's total reward for DSM activities in a given program year. The designs of shareholder incentive mechanisms include bonuses, rate-return adjustments, shared-savings, and hybrids combining two or more of these individual types of incentives (Eto et al. 1998). When program-specific incentives were not available, we chose to allocate a portion of total incentive payments to programs based on the energy saved by each program as a fraction of the total energy saved by all of the utility's DSM programs.

As discussed earlier, our decision to include shareholder incentives in our calculation of the total cost of saved energy may introduce an upward bias. An upper bound assessment of the magnitude of this bias is to assume that utilities do not bear any hidden costs and that all shareholder incentive payments are simply transfers. If we exclude shareholder incentives for the 27 programs that receive them, the simple average for total resource cost of saved energy of programs falls by about 7%.

ASSESSING UTILITY ESTIMATES OF ENERGY SAVINGS

Two quantities underlie utility estimates of energy savings from DSM programs: annual energy savings, and the economic lifetime of energy savings. As with the cost data provided, we relied on the information we received from the utilities on these quantities as final. However, following a review of the methods used by utilities to estimate savings, we conducted sensitivity analyses to determine the significance of potential biases.

Annual Energy Savings

All of the information on annual energy savings was based on some form of post-program savings verification. We classified methods for measuring annual energy savings into three broad categories: (1) tracking database methods, (2) billing analyses, and (3) end-use metering (Sonnenblick and Eto, 1995). Table 2 summarizes the annual energy savings methods used by the 40 programs. We also paid special attention to the treatment of spillover.

Most programs used more than one of the three methods to estimate annual energy savings. For example, all utilities maintain a tracking database of some sort to record information on program participants. Most utilities, however, augment their tracking databases to increase the reliability of their savings estimates. For example, the statistically adjusted engineering or SAE method reconciles a preliminary estimate of savings from a program's tracking database through a regression on customers' bills. Similarly, end-use metering is often used to refine estimates of hours of operation and, in some cases, changes in connected load. Thus, all of the various methods start from a tracking database.

Table 2. Summary of Annual Energy Savings Methods[a]

Tracking Database Methods	Number of Programs (% of 40)
Verification of Measure Installation	
On-Site - Sample	17 (43%)
On-Site - All	18 (45%)
Self-Report - Sample	4 (10%)
Self-Report - All	2 (5%)
Hours of Operation	
On-Site - Sample	17 (43%)
On-Site - All	3 (8%)
Self-Report - Sample	9 (23%)
Self-Report - All	9 (23%)
Based on Previous Study	6 (15%)
Billing Analyses	
Bill Comparison	1 (3%)
Bill Comparison w/Comparison Group	3 (8%)
Bill Regression w/Comparison Group	3 (8%)
SAE Regression w/Comparison Group	12 (30%)
End-Use Metering	11 (28%)

[a] For 11 programs, some methods were applied to only a subset of the energy saved by a program. For five programs, more than one method was used simultaneously to estimate savings for the program. The "number of programs" could add up to more than 40 because some programs used more than one method to verify installations or assess hours of operation.

Tracking Database Methods. Tracking database methods are often referred to as engineering estimates; however, we feel that this name is inaccurate because almost all evaluation methods involve some amount of engineering, so the name should not be applied only to tracking database methods. In addition, the word "engineering" obscures the fact that substantial post-program evaluation information is often incorporated into the estimate. This information ranges from the simple verification of program installations to detailed end-use metering of affected electrical circuits.

In its simplest form, the basic tracking database equation for annual energy savings consists of three terms:

Annual energy savings = Number of measures installed * Per measure changes in connected load * Hours of use

We distinguish among tracking database methods by the way in which additional information is introduced into this equation. Starting with the first term, measure installations can be verified by either on-site inspections (as in 35 of our programs) or customer self-reports (six programs).[5] On-site inspections may be conducted by utility staff, contractors to the utility, or both. Customer self-reports include information reported to the utility on a rebate application or through responses to telephone or mail surveys administered by the utility. The methods were applied either to all participating customers (21 programs) or to a sample of them (20 programs).

Changes in connected load are typically read from engineering tables that compare the connected load of the removed or replaced equipment to that of the more efficient replacement equipment. However, several utilities relied on end-use and short-term (described below) spot metering to verify and update these estimates.

Hours of operation were estimated through on-site inspections (20 programs), customer self-reports (18 programs),[6] and to a lesser extent on end-use metering (14 programs).[7] On-site inspections generally consisted of on-site interviews of customers. They were sometimes augmented by inspections of the premises to collect operating information for different zones within a premise. All (26 programs) or a sample (12 programs) of participating customers were surveyed.

For six programs, hours of operation were determined through tables that list "standard" hours of operation for specific end uses (such as lighting or chillers) usually with separate entries for different commercial building types (such as offices or schools). For five of these programs, the estimates were later either augmented by end-use metering or superseded by an SAE billing analysis. Use of these latter two methods decreases but does not eliminate concerns regarding the error that could be introduced by relying on look-up tables.

Billing Analyses. Billing analyses, in contrast to "bottom-up" tracking database methods, are a "top-down" approach for estimating savings. They are based, at a minimum, on monthly or annual billing information from participating customers, collected both prior to and after the installation of DSM

5. The numbers in this paragraph sum to more than the number of programs we examined because one utility employed more than one method for its program.

6. The number of programs relying on either self-reports or on-site inspections sums to less than the total number of programs because, as described in the next paragraph, some programs relied on look-up tables, which were later reconciled to actual bills using statistical techniques.

7. The numbers in this paragraph also sum to more than the number of programs we examined because utilities employed more than one method for some programs.

measures. Billing information can be analyzed using a simple differencing approach that directly compares pre-program to post-program consumption, with the individual bills sometimes first weather-normalized (as in four programs); they can also be analyzed using multivariate regressions (15 programs). The accuracy and reliability of estimates of net savings can be improved by including billing information from a comparison group of nonparticipating customers (done in 19 programs). A recent, very popular class of regression methods, called the statistically adjusted engineering or SAE method, relies on a preliminary estimate of savings (used by 12 programs). The coefficient emerging from the SAE regression measures is interpreted as a measure of the percentage of previously estimated savings that the regression model is able to confirm.

End-Use Metering. End-use metering is often regarded as the most accurate savings evaluation method because it measures the quantities most directly related to energy savings. However, because the cost of data collection is high, it is only used for only a small sample of participating customers. For the 14 programs that relied on end-use metering, between >1% and 12% of participating customers were metered. In absolute numbers, nine programs metered fewer than 40 customers, and two metered more than 50 customers. All of the metering studies are classified as short-duration studies, in which the metering periods generally last from two to four weeks.

Assessing Uncertainties in the Measurement of Annual Energy Savings. There are no generally accepted methods for measuring annual energy savings. All methods are subject to bias and imprecision. There is anecdotal evidence that the simplest forms of tracking database estimates of savings are biased upwards (Nadel and Keating, 1991). There is also some evidence to suggest that the realization rate determined using SAE models, which reconcile tracking database estimates to actual changes in energy bills, may be biased downwards (Sonnenblick and Eto, 1995). However, there is little information to judge bias and precision independently.

We conducted a preliminary examination of our 40 programs to see whether the methods used to estimate annual savings were systematically related to the resulting cost of saved energy. Among the 24 programs in which lighting accounted for more than 60% of savings,[8] we compared the measure cost component of the total resource cost of energy savings for the 15 programs that relied either on billing analyses or end-use metering to estimate savings to the

8. We made an assumption that savings due to lighting measures would be roughly identical since the programs tended to install the same types of lighting technologies (i.e., high-efficiency fluorescent lamps and ballasts) in similar situations (i.e., commercial buildings).

nine programs that relied solely on a tracking database to estimate savings.[9] The mean measure cost of the programs with savings based on tracking databases is slightly lower than the mean for programs with savings based on either billing analyses or end-use metering. Nevertheless, the standard deviations of the two means overwhelm the modest differences between them. Our data, therefore, do not support the existence of a statistically significant correlation between measurement method and annual energy savings.

Although this simple examination is by no means definitive, it suggests that savings are affected by other factors that have a greater impact than the type of savings evaluation methods used by a utility. This should come as no surprise since, as described, tracking databases vary greatly in the degree and quality of information they incorporate on actual installations. We conclude, in particular, that simple adjustments, such as the application of realization rates developed for one program to adjust the savings from another (used in Eto et al. 1996), cannot be justified without a more detailed understanding of the evaluation methods involved and the populations to which they were applied.[10]

Participant and Non-participant Spillover. Participant and non-participant spillover are savings that are caused indirectly by a utility's DSM program. Participant spillover refers to energy saving actions taken by a participant in addition to those supported by a utility's DSM programs (e.g., installing the same or different measures without a rebate). Non-participant spillover refers to energy savings actions taken by non-participants (i.e., those receiving no rebate) as a result of the program (e.g., by word-of-mouth or as a result of changes to industry standard practices). Spillover effects are important for examining the extent to which programs may be overcoming consumer reluctance to the adoption of energy-efficient measures and practices (Levine and Sonnenblick, 1994).

Evaluation methods for measuring spillover are in their infancy. Only two utilities made an explicit attempt to incorporate spillover in their estimates of program savings. Evaluations for 14 programs included survey questions on the subject of spillover. In several of these, the survey results were used to develop estimates or ranges of spillover savings. However, these spillover savings estimates were not included in the savings reported by the utility. Thus, to the extent that there are spillover effects from the programs, they are not accounted for and thus bias the total resource cost of saved energy upwards.

9. We focussed on the measure cost element (note: the measure cost element is expressed in ¢/kWh, not absolute $) of the total cost of saved energy because, to the extent that non-measure costs vary less than the cost of the measures with respect to quantity of energy saved, measure cost per unit of energy saved provides greater resolution on the impact of savings evaluation methods since savings are directly related to the number and hence cost of the measures installed.

10. See Sonnenblick and Eto (1995) for a longer discussion of the challenges that must be addressed by the practice.

Estimating the Economic Lifetime of Savings

The economic lifetime of savings is the second element required to establish cumulative energy savings. The estimation of the economic lifetime of savings remains a critical source of uncertainty in the measurement of energy savings from utility DSM programs. It will be several years before it is possible to conduct definitive studies to determine the long-term persistence and economic lifetime of savings from many of the most popular DSM measures because many DSM technologies are new to the market. More commonly, utilities have conducted short-term persistence studies to determine measure retention, removal, and failure for periods of one to four years following installation (Wolfe et al. 1995).

Information on the lifetime of savings was generally reported separately for each measure or as a savings-weighted aggregate for all measures (33 programs). We did not receive information on measure or savings lifetimes for seven programs. We developed estimates for three of these programs by constructing a weighted average based on the largest contributors (weighted by either savings, measures, or participants) to savings. For two programs, in which savings were not reported by measure or participant, we made an estimate based on lifetimes reported for programs offering similar measures. For two programs, we used the lifetimes for the popular measures installed. Measure lifetimes range from six to 18 years. The simple average is 13.0 years with a standard deviation of 3.1 years.

In our sample, eight programs had completed measure persistence studies that included the 1992 program year. (Typically, the studies include other program years as well.) The studies generally found high rates of persistence for most measures. Notably, several of these studies found low renovation rates after installation of DSM measures in offices, restaurants, and retail premises in contrast to earlier, well-reported findings of high (25% or more per year) renovation rates in these types of premises (Hickman and Brandeis, 1992). Renovation is an important consideration for DSM measure lifetimes because renovation typically involves replacement of (in this case, energy-efficient) equipment prior to the end of its physical lifetime.

Assessing Uncertainties in the Estimation of Economic Lifetimes of Savings It is straightforward to calculate the sensitivity of the total resource cost of energy savings to different savings lifetimes. For a program with a savings lifetime of 13 years and a total resource cost of energy savings of 4¢/kWh, a decrease in savings lifetime to 10 years increases the cost by 22%, and an increase in savings lifetime to 16 years decreases the cost by 13%.

We also examined the sensitivity of our findings to the economic lifetime of savings assumed by the utilities by replacing reported lifetimes with a standard set of assumptions. For programs in which lighting savings accounted for more than 60% of savings, we assumed a lifetime of 10 years (24 programs). For the remaining programs (in which lighting accounted for less than 60% of savings), we assumed a lifetime of 14 years (12 programs).

We found that the use of standard measure lifetimes increases the mean total resource cost of saved energy by about 10%, but that the increase is not statistically significant. In particular, use of standard lifetime estimates does not reduce variance in total resource costs. We conclude that uncertainty in the total resource cost of saved energy resulting from reliance on lifetimes that are estimated (out of necessity) is not significantly reduced through the use of standard assumptions.

THE TOTAL RESOURCE COST OF SAVED ENERGY AND PROGRAM COST-EFFECTIVENESS

We find that, on a savings-weighted basis, the total resource cost of energy saved by the 40 programs is 3.2 ¢/kWh (see Table 3).[11] Measure costs. split between utility and participants, account for 44% (1.4 ¢/kWh) and 31% (1.0 ¢/kWh), respectively. The large fraction accounted for by customer-paid measure costs (31%) highlights the importance of including these costs in a full accounting of the total cost of energy savings to society. Ignoring these costs would make the apparent cost of energy savings one third less expensive than they actually are. Utility non-measure costs, which include utility overhead, evaluation, and shareholder incentives, account for 0.8 ¢/kWh or 25%.

Figure 2 arranges the 40 DSM programs from the least to the most expensive and plots them sequentially against energy savings; the "width" of each program along the x-axis represents the annual energy savings accounted for by each program. This presentation shows that the savings-weighted average is dominated by several very large and inexpensive programs, and that the most expensive programs were comparatively small in size. For example, 28% of the savings have cost less than 2 ¢/kWh and 50% have cost less than 3 ¢/kWh. At the same time, only 1% have cost more than 9 ¢/kWh.

The savings-weighted total resource cost of energy savings (3.2 ¢/kWh) is almost 20% lower than previously reported findings for 20 commercial lighting programs, which presented a savings-weighted total resource cost of energy savings of 3.9 ¢/kWh (Eto et al. 1996). Notably, the earlier findings did not include shareholder incentives. We believe there are two reasons for the difference. First, the results for our sample are strongly affected by the presence

11. Following Joskow and Marron (1992), the total resource cost of saved energy is calculated by dividing the levelized cost of a program by annual energy savings. Hence, the TRC is expressed as a cost per kilowatt-hour of savings (¢/kWh). All levelizations were performed using a common real (i.e., net of inflation) discount rate of 5%.

Table 3. The Total Resource Cost of Saved Energy (¢/kWh)

Program ID	Measure Costs		Non-measure Costs		Total Resource Cost of Saved Energy	Avoided Costs	TRC Ratio
	Utility-paid	Participant-paid	Administration and Evaluation	Shareholder Incentives			
1	1.9	1.7	0.4	0.7	4.7	6.6	1.4
2	1.8	0.0	3.2	0.6	5.6	11.2	2.0
3	1.9	0.6	0.3	0.2	3.0	8.2	2.7
4	3.5	0.0	0.3	0.8	4.6	4.0	0.9
5	1.0	4.1	0.4	0.5	5.9	5.3	0.9
6	1.7	6.7	0.4	0.6	9.3	5.5	0.6
7	1.6	1.5	0.1	1.6	4.8	6.6	1.4
8	7.5	2.0	0.9	0.3	10.7	6.6	0.6
9	33.8	2.9	21.8	(10.4)	48.1	6.6	0.1
10	1.1	1.3	0.5	0.1	3.0	3.1	1.0
11	1.0	0.0	1.6	0.1	2.7	3.7	1.4
12	1.1	0.2	0.1	0.4	1.7	5.1	2.9
13	1.0	3.8	0.6	0.1	5.5	5.2	1.0
14	1.3	1.7	0.2	0.0	3.2	3.0	1.0
15	1.9	1.3	0.2	0.0	3.4	4.0	1.2
16	1.7	3.4	0.4	0.0	5.5	8.9	1.6
17	5.5	0.0	1.2	0.1	6.8	10.7	1.6
18	12.5	1.2	3.9	0.1	17.6	9.8	0.6
19	2.6	0.0	0.3	0.1	3.0	7.9	2.7
20	1.8	0.4	0.4	0.0	2.5	4.5	1.8
21	1.5	0.7	0.3	0.0	2.5	4.5	1.8
22	4.3	0.0	1.5	0.0	5.8	12.1	2.1
23	2.0	1.6	2.4	0.0	5.9	12.1	2.0
24	2.4	0.6	0.5	0.4	4.0	6.7	1.7
25	7.1	0.0	0.9	0.6	8.5	10.1	1.2

Table 3. Continued

Program ID	Measure Costs		Non-measure Costs		Total Resource Cost of Saved Energy	Avoided Costs	TRC Ratio
	Utility-paid	Participant-paid	Administration and Evaluation	Shareholder Incentives			
26	2.6	0.7	0.6	0.2	4.1	7.1	1.7
27	7.0	0.0	1.1	0.3	8.4	10.0	1.2
28	5.6	0.0	0.5	0.0	6.1	5.9	1.0
29	1.2	2.3	0.5	0.0	3.9	5.2	1.3
30	4.5	0.0	0.9	0.6	6.0	4.8	0.8
31	2.2	0.9	0.4	0.6	4.1	5.4	1.3
32	0.3	0.4	0.1	0.0	0.8	7.7	9.6
33	0.7	0.4	0.3	0.7	2.1	7.0	3.3
34	1.4	0.0	0.3	0.8	2.6	10.4	4.1
35	0.3	0.8	0.3	0.6	2.0	17.6	9.0
36	0.3	1.0	0.1	0.0	1.5	3.1	2.1
37	0.8	0.8	1.9	0.0	3.5	4.4	1.3
38	5.7	0.0	0.8	0.8	7.3	5.6	0.8
39	1.5	1.5	0.6	0.3	4.0	5.6	1.4
40	2.1	0.2	1.1	0.4	3.9	5.6	1.4
Weighted Average	1.4	1.0	0.4	0.4	3.2	6.6	3.2
Mean	3.5	1.1	1.3	0.1	6.0	6.9	1.9
Standard Deviation	5.5	1.4	3.4	1.7	7.5	3.1	1.9

of large, inexpensive programs, which, in contrast to the convenience sample examined in the previous study, were included by design in the current study. Second, information for the current study comes in most cases from 1992 program year data; the previous study used data from 1991 or earlier. For programs that appear in both studies, this means that the data have been updated for the current study. Several of these programs have reduced the cost of acquiring energy savings.

We describe program cost-effectiveness using a standard DSM benefit-cost test, the total resource cost (TRC) benefit-cost test ratio. In the total resource cost test, the value of DSM programs is measured by the resource costs they allow the utility to avoid (CPUC/CEC, 1987). Avoided resource costs depend on the economic circumstances of a particular utility and on the load shape impacts and economic lifetime of savings from a particular DSM program (Busch and Eto, 1996). Definitions of what cost components are avoided and what methods are used to estimate them differ among utilities. For example, we eliminated environmental externality adders not because we think they are unimportant but in an effort to ensure greater comparability among utilities.

The overall TRC benefit-cost test ratio of total avoided costs to total resource costs is 3.2, indicating that, taken as whole, the programs are highly cost effective. See Figure 3. The simple average of the TRC benefit-cost test ratios is 1.9 with a standard deviation of 1.9. Because the overall TRC benefit-cost test ratio is higher, we can conclude that some of the largest programs are also the most cost effective. The high standard deviation also indicates that some programs are not cost effective; 11 of the programs have TRC benefit-cost test ratios of less than 1.0. This should not be too surprising because there are several extremely high-cost programs. The 11 programs that are not cost effective account for 12% of the total resource costs of all of the programs.

The most critical issue for our estimates of program-specific avoided costs, and hence these cost-effectiveness results, is that they are based on a forecast of the future and hence are inherently uncertain. For many utilities, avoided costs have dropped significantly since the time when the programs were first developed. In particular, the program planning estimates for our 1992 programs were for the most part based on estimates of avoided cost developed in 1991. In view of this situation, it is useful to consider how lower avoided costs would affect our findings. If we assume that avoided costs are 50% lower than those originally reported, TRC benefit-cost test ratios drop below unity for an additional 19 programs. However, the overall TRC benefit-cost test ratio would be 1.6. We conclude that dramatically lower avoided costs can have a significant effect on the cost effectiveness of individual programs. Nevertheless, taken as a whole, the majority of savings from the programs remain cost effective.

Figure 2. The Total Resource Cost of Saved Energy from the Largest 1992 Commercial Sector DSM Programs

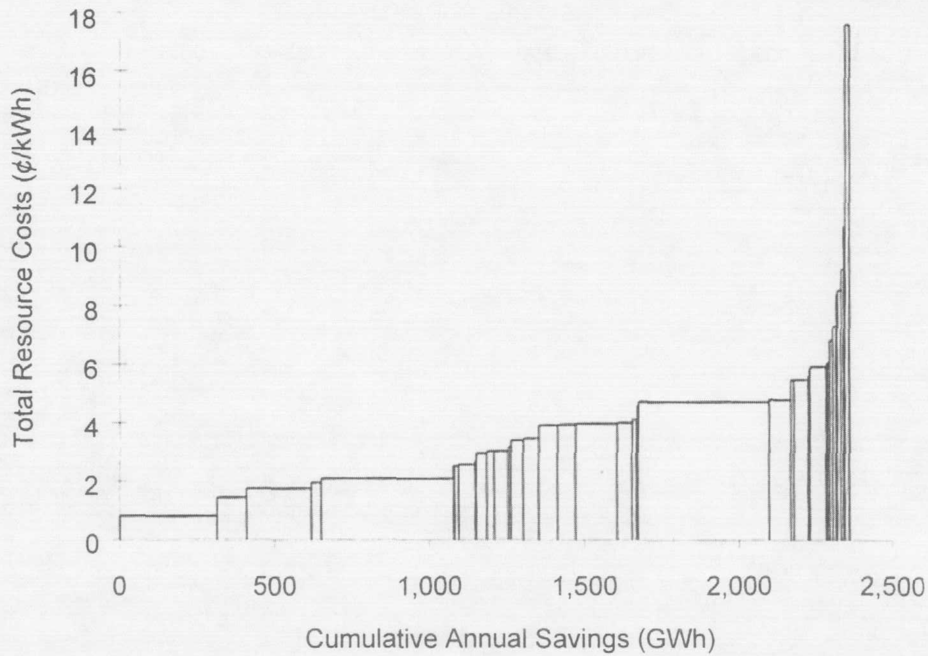

Note: Programs are arranged in ascending order from the least to most expensive (¢/kWh). The horizontal "width" of the programs represents annual energy savings. The savings-weighted mean is 3.2 ¢/KWh; the unweighted mean is 6.0 ¢/kWh.

Figure 3. The Cost-Effectiveness of the Largest 1992 Commercial Sector DSM Programs

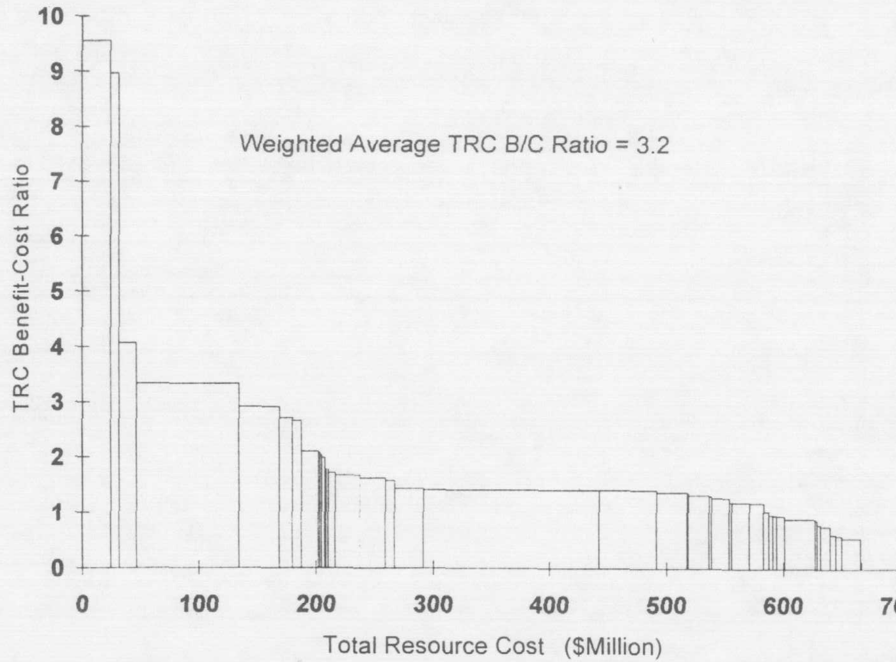

Note: Cost effectiveness is expressed by the ratio of avoided costs to the total resource cost of saved energy. Programs are arranged in descending order from the most to the least cost effective. The horizontal "width" of the programs represents the total resource cost. The overall ratio of total avoided costs to total resource costs is 3.2. *The simple average of individual benefit-cost ratios is 1.9.*

EXAMINING DIFFERENCES IN THE COST OF SAVED ENERGY

What makes some DSM programs more costly than others? We tried to address this question in a preliminary fashion by reviewing the differences among groupings of the least and the most expensive programs. We also conducted multiple regression analysis of the correlation between various program design and operational features and the cost of saved energy. As the regression results added little to our initial findings, we refer to the interested reader to a longer discussion of these results in a separate document (Eto et al. 1995).

Table 4 summarizes aspects of the five least and five most expensive programs as well as for the entire sample of programs. Starting with program costs, it is clear that the least expensive programs were run with substantially lower administrative or non-measure costs. These programs were either run more efficiently (even including, or perhaps because of, shareholder incentives) or they were able to spread fixed, non-measure costs over a larger base of energy savings (see below). It is also clear that high measure costs account for most of the costs of the more expensive programs. Either the measures installed were very costly or the installations were such that comparatively less energy was saved per installation (e.g., facilities receiving the measures had few hours of operation).

For the least expensive programs, several features stand out. First, the programs were very large, as measured either by annual savings or by number of participants. (See also Figure 2.) Second, these are some of the older, possibly more mature programs in the sample. Third, because the programs tended to report incremental measure costs, they appear to have targeted normal rather than early replacements.

The most expensive programs also have some common features. First, the programs are quite small as measured by total savings. Second, they appear to be somewhat newer programs compared to the entire sample. Both these factors suggest that these programs are not fully mature, so fixed administrative costs are being spread over a smaller base of savings. Finally, they include more direct-install programs, for which full measure costs would be reported.

Perhaps more interesting than the differences between the least and most expensive programs are the similarities between them. Avoided costs, the percentage of measure costs paid by the utility, lighting fraction of total savings, economic lifetime of savings, and savings evaluation methods are all quite similar to one another. These similarities have important implications for previous findings and for the potential impact of methodological differences on current findings.

Table 4. Comparison of Characteristics of the Most and Least Expensive Programs

	Five Least Expensive (¢/kWh) Programs	Five Most Expensive (¢/kWh) Programs[a]	All Programs[a]
Total Resource Cost of Saved Energy (¢/kWh)	1.6 (0.5)	10.9 (3.9)	4.9 (3.1)
Nonmeasure Cost (¢/kWh)	0.5 (0.4)	1.8 (1.2)	1.1 (0.9)
Shareholder Incentives	4 of 5	5 of 5	28 of 39
Measure Cost (¢/kWh)	1.1 (0.3)	9.1 (2.7)	3.8 (2.6)
Avoided Costs (¢/kWh)	8.1 (5.6)	8.4 (2.2)	6.9 (3.1)
Measure Costs Paid by Utility (%)	49 (26)	78 (34)	69 (28)
Program Size (GWh/yr)	215.3 (159.5)	8.7 (5.0)	60.4 (104.8)
Participants (per year)	4,721 (4,626)	796 (1,095)	1,691 (2,563)
Savings/Participant (MWh)	106.3 (120.1)	48.2 (80.9)	71.8 (124.8)
Lighting Fraction of Total Savings (%)	53 (33)	73 (42)	64 (39)[b]
Program Type	0 of 5 Direct Install	3 of 5 Direct Install	10 of 39 Direct Install
Program Start Date	1987 (5)	1990 (0)	1989 (3)
Economic Lifetime of Savings (Years)	14.4 (1.2)	11.3 (2.2)	13.1 (3.1)
Savings Evaluation Method (Billing, Metering, Tracking)	3 of 5 Bill or Meter	3 of 5 Bill or Meter	24 of 39 Bill or Meter

Note: Where means are presented, standard deviations are also reported in parentheses.
[a]This analysis does not include one very high-cost program, which was greater than two standard deviations above the mean.
[b]N = 37 for this explanatory variable.

Eto et al. (1996) found evidence suggesting that avoided costs were positively correlated with total resource cost of saved energy and concluded that avoided costs helped to explain the differences in program costs. They speculated that avoided costs could be thought of as the value standard against which utilities designed programs. In this situation, higher avoided costs led to higher-cost programs. In the current situation, the explanation appears more complicated, probably because of confounding influences, such as program size, type, and maturity.

The similarity in the portion of savings attributable to lighting and the similarity in the lighting measures promoted suggests that differences in the portfolio of technologies offered by a program may have been less important than the savings that resulted from the specific technologies that were actually installed and the use of these technologies in the premises in which they were installed. Earlier we speculated that the more expensive programs may have ended up targeting installations with lower savings (because of a small number of hours of operation, for example). This similarity in lighting savings fractions lends some credence to this hypothesis. A definitive conclusion can only be drawn by examining detailed demographic information on actual installations.

CONCLUSION

We conducted a detailed review of the largest commercial sector DSM programs operated by electric utilities in 1992 to determine the total cost to society of the energy they saved. We documented the significance of consistently including information on: 1) customer contributions to cost of energy-efficiency measures, 2) the baseline against which the incremental cost of energy efficiency is calculated, and 3) non-measure costs, including administration, evaluation, and the incentives paid to utilities in return for operating DSM programs effectively. We also examined current methods for evaluating annual savings and found that the choice of method did not have a statistically significant impact on the cost of saved energy. At the same time, we confirmed that estimates regarding the economic lifetime of savings do have a significant influence on the cost of saved energy. We considered individual sources of biases that we may have introduced in order to develop consistent cost and savings information and concluded that, to the extent we could determine we might have introduced them, doing so led to overstatement of the cost of saved energy. We found that the programs, as a whole, had saved energy at a cost of 3.2 ¢/kWh and, compared to the cost of the energy they allowed the utilities to avoid generating or purchasing (in the absence of these programs), that they were cost effective. We were then able to relate selected program design and operational features to the cost of saved energy.

No one knows the future of utility ratepayer-funded energy-efficiency programs. In some parts of the US, strong public support for these and other public purpose activities (such as low-income programs, renewable energy development, and research and development) has led to the creation of wires charges to continue funding for these programs in a restructured electricity industry (Eto, Goldman, Nadel 1998). In addition to the issue of whether funding should be continued, the role of utilities in delivering energy-efficiency programs (as opposed to simply collecting the funds through rates for programs that are administered by others) remains an important subject of discussion. We hope that the information presented in this paper can inform these and other discussions regarding the wisdom of public funding for large-scale efforts to promote energy efficiency that rely primarily on financial incentives and targeted information to accelerate voluntary adoption of energy-efficient measures.

Acknowledgments

The work described in this paper was funded by the Assistant Secretary of Energy Efficiency and Renewable Energy, Office of Utility Technologies, Office of Energy Management Division of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098. The Database on Energy Efficiency Programs project, which was the source of information for this paper, has also received funding from the New York State Energy Research and Development Authority, the Bonneville Power Administration, the Rockefeller Family and Associates, the Electric Power Research Institute, and in-kind support from the Energy Foundation.

This paper could not have been prepared without the information provided by and cooperation of the following: Baltimore Gas & Electric, Boston Edison, California Public Utilities Commission, Central Hudson Electric and Gas, Central Maine Power, Connecticut Public Utilities Commission, Consolidated Edison, Consumers Power, Green Mountain Power, Long Island Lighting, Maine Public Utilities Commission, Michigan Public Service Commission, Montana Power, New England Electric System, New York State Department of Public Service, New York State Electric and Gas, Niagara Mohawk Power, Orange and Rockland, Pacific Gas and Electric, Potomac Electric Power, Puget Sound Power and Light, Rochester Gas and Electric, Sacramento Municipal Utility District, San Diego Gas and Electric, Southern California Edison, Texas Public Utilities Commission, Texas Utilities Electric, and Western Massachusetts Electric. We gratefully acknowledge the support of these firms and institutions.

REFERENCES

Busch, J. and J. Eto (1996). "Estimation of Avoided Costs for Electric Utility Demand-Side Planning." *Energy Resources* 18: 473-499.

California Public Utilities Commission and California Energy Commission (CPUC/CEC) (1987). "Standard Practice Manual, Economic Analysis of Demand-Side Management Programs." P400-87-006. Sacramento CA. December.

Energy Information Administration (EIA) (1998). "Electric Power Annual, Volume II." DOE/EIA-0348(97)/2. Washington, D.C. October.

Eto, J., C. Goldman and S. Nadel (1998). "Ratepayer Funding for Energy Efficiency in a Restructured Electricity Industry: Issues and Options for Regulators and Legislators." LBL-41479. Lawrence Berkeley National Laboratory, Berkeley CA, May. http://eetd.lbl.gov/EA/EMP

Eto, J., S. Kito, L. Shown, and R. Sonnenblick (1995). "Where Did the Money Go? The Cost and Performance of the Largest Commercial Sector DSM Programs." LBL-38201. Lawrence Berkeley National Laboratory, Berkeley CA, December. http://eetd.lbl.gov/EA/EMP

Eto, J., S. Stoft, S., and S. Kito (1998). "DSM Shareholder Incentives: Current Designs and Economic Theory." *Utilities Policy* 7: 47-62.

Eto, J., E. Vine, L. Shown, R. Sonnenblick, and C. Payne (1996). "The Cost and Measured Performance of Utility-Sponsored Energy Efficiency Programs." *The Energy Journal* 17 (1): 31-51.

Hickman, C. and P. Brandeis (1992). "Building Site Visits, A Look at Measure Retention as an Element of Program Evaluation." Proceedings. ACEEE 1992 Summer Study on Energy Efficiency in Buildings. American Council for an Energy Efficient Economy, Washington, DC.

Joskow, P. and D. Marron (1992). "What Does a Negawatt Really Cost? Evidence from Utility Conservation Programs." *The Energy Journal* 13 (4): 41-74.

Levine, M. and R. Sonnenblick (1994). "On the Assessment of Utility Demand-Side Management Programs." *Energy Policy* 22 (10): 848-856.

Nadel, S. and K. Keating (1991). "Engineering Estimates versus Impact Evaluation Results: How Do They Compare and Why?" Proceedings. 1991 International Energy Program Evaluation Conference. Chicago IL.

Parfomak, P. and L. Lave (1996). "How Many Kilowatts are in a Negawatt? Verifying Ex Post Estimates of Utility Conservation Impacts at the Regional Level." *The Energy Journal* 17 (4): 59-87.

Sonnenblick, R. and J. Eto (1995). "A Framework for Improving the Cost Effectiveness of DSM Program Evaluations." LBL-37158. Lawrence Berkeley National Laboratory, Berkeley CA. http://eetd.lbl.gov/EA/EMP

Wolfe, A., M. Brown, and D. Trumble (1995). "Measuring Persistence: A Literature Review Focusing on Methodological Issues." ORNL/CON-401. Oak Ridge National Laboratory, Oak Ridge TN.