

Actuarial Pricing of Energy Efficiency Projects: Lessons Foul and Fair

Paul Mathew, Ph.D., Lawrence Berkeley National Laboratory,
901 D. Street SW, Suite 950, Washington DC 20024.
phone: 202-646-7952, fax: 202-646-7800

Steven Kromer, PE, Teton Energy Partners
1639 Delaware Street, Berkeley, CA 94703

Osman Sezgen, Lawrence Berkeley National Laboratory
1 Cyclotron Rd., Mail Stop 90-4000, Berkeley, CA 94720

Steven Meyers, Rational Energy Network
4500 Spanish Oak Trail, Austin, TX 78731

Abstract

Recent market convulsions in the energy industry have generated a plethora of post-mortem analyses on a wide range of issues, including accounting rules, corporate governance, commodity markets, and energy policy. While most of these analyses have focused on business practices related to wholesale energy trading, there has been limited analysis of retail energy services, particularly energy efficiency projects. We suggest that there were several business concepts and strategies in the energy efficiency arena whose inherent value may have been masked by the larger failure of companies such as Enron. In this paper, we describe one such concept, namely, actuarial pricing of energy efficiency projects, which leverages a portfolio-based approach to risk management. First, we discuss the business drivers, contrasting this approach with conventional industry practice. We then describe the implementation of this approach, including an actuarial database, pricing curves, and a pricing process compatible with commodity pricing. We conclude with a discussion of the prospects and barriers for the further development of risk management products for energy efficiency, and the development of energy efficiency as a tradeable commodity. We address these issues from an experiential standpoint, drawing mostly on our experience in developing and implementing such strategies at Enron.

Keywords: energy efficiency, actuarial pricing, risk management

1 Introduction

The deregulation of electric power industry, particularly since the mid-1990s, created enormous new business opportunities for the development of new risk management products and services (Sioshansi 2002, Chamberlin and Herman 1995). Companies such as Enron, whose core expertise was in energy assets and risk management, capitalized on this opportunity and grew dramatically. Having successfully transformed the wholesale electric power and natural gas markets, Enron sought to expand into the retail energy business through its newly established business unit Enron Energy Services (EES). EES' primary focus was to leverage its expertise in risk management to sign commodity deals (i.e. contracts for the supply of electricity and natural gas) with relatively large retail customers (annual energy expenditures greater than \$10 million). Enron soon realized that there was also a significant business opportunity for "bundling" commodity contracts with energy asset related services such as energy efficiency projects. Indeed, in some cases, customers requested these services, as part of an overall out-sourcing of energy asset management. In order to meet customer demand for energy efficiency services, EES created a group dedicated to optimizing the costs of owning and operating energy-using assets.

While the concept of assisting facility owners find value in their infrastructure was not new, the task of identifying, implementing and verifying energy projects on this massive scale (about \$350 million in capital approved in two years) required new approaches. The challenges for EES were to reduce the transaction costs associated with implementing efficiency projects and create a knowledge base to inform future efficiency investments. EES looked to other financial/risk management industries and found a potential solution – actuarial pricing. This paper describes the challenges, successes and problems encountered while developing and deploying this new approach to energy efficiency.

2 Business drivers for actuarial pricing

While the need for more cost-effective systems to deploy and manage energy efficiency projects should be self-explanatory, the method proposed by Enron - actuarial pricing - requires an understanding of the business context. The primary area of innovation throughout Enron was in taking a new look at old industries and their business processes and rewriting those processes through the lens of standardization, commodification and tradability. These were the backbone of commodity markets, made all the more effective by harnessing the internet. And since Enron's first big successes with this approach had been in the wholesale markets for natural gas and electric power, retail electric and gas contracts seemed the logical next step.

The value of energy efficiency is the costs that are avoided by providing services with less energy. The avoided costs are directly related to the amount of energy that is saved. Enron had much experience and powerful systems to predict the price of energy, but was lacking in systems to predict the amount that would be saved at a particular cost in a particular market. What was missing was a standard approach to quantifying the results of events in the physical world – in effect to commodify energy services.

There is, however, a dissonance between the commodity trading model and energy efficiency business as it is typically practiced. Energy efficiency projects are, ultimately, engineering projects involving real assets, while commodity contracts are, in essence, financial transactions. Consider the way an energy efficiency project is typically priced. The Energy Service Company (ESCO) conducts one or more engineering audits of a customer's energy assets to assess the opportunities for energy savings projects. Engineering calculations of the potential energy savings are then made using the data collected in the audits. Audits and engineering calculations are, in effect, the transaction costs for energy savings projects. These costs can significantly impact the financial viability of a project – especially small projects, where the cost of an audit could easily exceed a year's worth of energy savings.

This approach runs counter to commodity pricing, in three important ways:

- The pricing cycle time was too long – technical personnel had to visit dozens of sites, gather data, and then perform relatively labor-intensive calculations. As a consequence it was often the case that the pricing of the energy savings projects delayed the pricing of the overall bundled contract.
- It precluded a “low-touch high-volume” sales process. Scalability was a critical aspect of Enron’s business goals.
- It was not consistent with the business systems and processes (e.g. information technology infrastructure) for commodity contracts.

These factors essentially forced Enron to reconsider the way energy savings projects were priced, and initiated the development of a quasi-actuarial approach to pricing energy efficiency projects. Essentially, it drew an analogy with the pricing of individual insurance policies. The analogy worked thus: If a customer seeks to get auto insurance, she can call up an insurance provider, who then asks her five to ten questions (location, type of car, age and gender of driver, etc.) and provides a quote based on her responses. The provider is able to do this because it has actuarial tables developed from statistical models that relate the probability of future claims to the pertinent customer characteristics. Thus, the provider does not need to ask the customer for a driving test, medical test, etc. in order to determine her “risk profile”.

Could energy efficiency projects be priced in a similar manner? That is, could the overall value of energy savings projects in a customer portfolio be assessed using pertinent customer site characteristics, without doing detailed engineering audits on individual sites? Such an approach could significantly reduce transaction costs, and thereby broaden the scope of viable projects.

Enron aggressively sought to answer these questions by actually trying to implement actuarial pricing. In the next section, we describe the development, implementation and use of an actuarial pricing process for energy efficiency projects.

3 Actuarial pricing tools and process

3.1 Actuarial database of energy efficiency projects

The fundamental prerequisite for an actuarial approach is data from which to develop probabilistic predictive models. In the case of the automobile insurance pricing, for example, the providers have large databases of automobile accident data, insurance claims, and customer characteristics.

Therefore, given the charge to develop an actuarial approach to value energy savings projects, the first task was to develop a database of energy savings projects. This database would then be used to develop “pricing curves” for energy efficiency, analogous to the forward curves used in commodity pricing. We describe curve development and use in more detail later in section 3.2.

While the theoretical basis for such an approach has precedent, and the ESCO industry has been implementing energy efficiency projects for over a decade, there are nonetheless several challenges in developing such an actuarial database:

- Lack of documentation: Many ESCOs simply do not document and archive their project data in a manner that is effectively retrievable for analysis.
- Lack of standardization: There is no standard or systematic way of documenting and reporting the results of projects, even within one organization.
- Proprietary data: Most ESCOs do not share the data from their projects, since this is viewed as competitive information. However, data from utility DSM programs were available from several utilities, after removing identifying information.

Thus, while there was in principle a significant amount of data “out there,” their effectiveness for actuarial use was limited for the reasons cited above. One notable exception was the DOE Industrial Assessment Center (IAC) Database (OIT 2003). This government-funded, public domain database has data in standardized format for over 60000 industrial energy efficiency projects, although not all were implemented. (It’s worth noting that the IAC database is a good example of public-private partnerships to develop shared information resources to advance market transformation in energy efficiency.) The National Association of ESCOs (NAESCO) maintains a proprietary database of about 800 self-selected ESCO projects. This database has been used to analyze trends in the ESCO industry (Goldman et al. 2001), but does not have enough detail on each project (system parameters, etc.) to conduct actuarial analysis.

Clearly, there is no single “correct” data schema for such a database – rather, the schema and the pricing methodologies will have to evolve based on empirical considerations rather than be analytically derived from a theoretical basis. Based on our experience at Enron, we briefly describe below the major elements of a data schema for energy efficiency projects.

ECM Type: There should be a list of standard ECM (energy conservation measure) types, hierarchically categorized by technology area (Figure 1). The major challenge here is to find the right balance in terms of granularity. A list that is too fine-grained may result in too many ECMs being classified as “combination,” while one that is too coarse will limit the actuarial significance of ECM type.

System		ECM Group	ECM Type		
1	Lighting	g31 Upgrade	311	Dry bulb economizer	
2	Building Envelope		312	Enthalpy economizer	
3	Packaged Units		313	Heat recovery media	
4	Air Distribution		314	Clean and charge	
5	Boilers		315	Replace w/ high efficiency unit	
6	Chillers		316	Convert dump damper to VFD	
7	Motors		317	Rebalance and reduce air flow	
8	Refrigeration		318	Liquid pressure amplifier	
9	Compressed Air		319	Electronic expansion valve	
10	Building Automation		320	Refrigerant subcooler	
11	Electrical Systems		321	Evap condenser retrofit	
12	Generation		322	Refrigerant oil additives	
13	Maintenance Measures		323	Retrofit air cooled to water cooled	
14	Industrial Processes		g32 Controls	321	Night setback/setup
15	Water Treatment			322	Schedule change
xx	...	323		CO2 ventilation control	
		g38 Combination	398	Combination	

Figure 1: Hierarchical categorization of ECM types for actuarial database

ECM parameters: For each ECM type there should be a standard set of equipment specifications and operational parameters, both baseline and post-retrofit. Obviously, a more exhaustive list of parameters will allow for a more robust analysis of parameter significance. However, actual experience with data collection suggest that if the data input requirements are seen as too onerous, project developers are less likely to take the effort to input data into the database. Figure 2 indicates a sample of ECM parameters for packaged rooftop units.

Name	Description	Unit	Values
PUAge	Average age of the units	years	<5, >5 but <15, >15
PUSetback	Setback controls	-	Yes, No, N/A
PUCtrlType	Control type	-	EMS, Timeclock and Tstat
PUDeConVent	Demand controlled ventilation	-	Yes, No, N/A
PUEcon	Economizer type	-	Temperature, Enthalpy, N/A
PULastMaint	Year since last maintenance	years	<1, >1 but <3, >3
PUNum	Total number of units	-	
PURunHrs	Average annual operating hours of units	hours	
PUTons	Total capacity of all units	tons	

Figure 2: Limited set of actuarial parameters for packaged roof top units

Energy Savings: In addition to the standard data pertaining to energy baseline and savings for different fuels, the data schema should include uncertainty parameters for each quantity. At a minimum these include mean, p5, p95, and the associated measurement and verification (M&V) method for each quantity. This will allow for actuarial analysis of the “optimum” amount of M&V based on risk tolerances.

Implementation cost: This should preferably be broken out into materials, labor volume and time, audit & design costs, taxes, overhead & profit. However, in some cases (e.g. lump sum contracts), contractors are hesitant to provide this breakdown, since it is considered competitive information.

Schedule: The time taken to complete the project affects the present value of savings. It should preferably be broken out into design, construction, and commissioning time.

An actuarial database with this schema was implemented at Enron using MySQL, with an MS Access front end. The database had almost 100,000 records of project data. The next two sections describe how this database was used to develop curves for actuarial pricing.

3.2 Pricing curve development and risk management

Figure 3 illustrates an ECM savings curve developed from the actuarial database. This particular curve describes the annual electrical savings from using setback controls for roof top packaged HVAC units (RTUs) in office buildings, expressed in annual kWh savings per ton RTU capacity. The curve is further specific to a particular region (northeast US) and equipment age (more than 5 yrs). This curve is essentially a histogram of the savings from similar projects that are recorded in the database – including both pre-construction estimates and post-construction “measured” savings. The range of savings is shown as a standard and cumulative probability distribution. The numerical percentile values are shown on the right. Curves are similarly developed for ECM project cost and schedule.




Figure 3: Sample ECM savings curve generated from actuarial database.

Note that for a given metric, curves can be developed at various levels of specificity, as conceptually shown in figure 4. Thus, in theory, the more site-specific parameters one is able to obtain, the more context-specific the curve will be. However, there is a tradeoff between the specificity of the curve and the number of data points it is generated from.




Figure 4: Conceptual illustration of savings curves at different levels of specificity.

Therefore, the use of this curve in a pricing process requires an assessment of its “quality”, taking into account parameters such as the number of data points, whether the data was measured or estimated, etc. For instance, a curve will be of higher quality if it is based on a large number of data points and if those data points are based on measurements rather than estimations. Enron developed a process to “score” each curve.

It’s clear that the quality score of a price curve needs to be quantitatively derived from the risks associated with the curve. Estimating uncertainty that is specific to energy savings projects has been described in various publications (ASHRAE 2002, Jump 2001, Reddy et al. 1999), although it is still in its infancy. Moreover, most of the work to date does not address the treatment of uncertainty at the portfolio level. For a large portfolio, a distinction needs to be made between risks that can be diversified across the portfolio (diversifiable risk) and those which cannot (non-diversifiable risk). The non-diversifiable risk is typically charged to the customer as a risk premium, and is therefore the basis for deriving the curve score.

Consider again a savings curve that is derived from a database, using both estimated and measured data from similar projects. In an ideal world, a large database of measured projects would capture the inherent risk of true savings. In reality, the paucity of available measured project data results in a database that has a large number of estimated values and a small number of measured values. Therefore, the data used to develop a savings curve need to be corrected for estimation bias.

Consequently, there are two main sources of risk: First, there is an inherent uncertainty of the savings-estimate, due to various unknown or unknowable factors that affect the actual savings (e.g. exact schedule of occupancy and operation). Secondly, there are inaccuracies in the way the project developers estimate the savings — the measured-deviation, which is determined by calculating the difference between the estimated values and the measured savings for similar projects. The mean of the distribution of measured-deviation is the estimation bias.

The *savings-estimate* and *measured-deviation* each exhibit a distribution with a mean and standard deviation. In a large portfolio of projects, the risks associated with the uncertainty of savings-estimate as well as measured-deviation will be diversified and the result for the portfolio will approach the mean values for each of them. However, the uncertainty about the accuracy of the mean values themselves is a non-diversifiable risk. In other words, no matter how big the portfolio is, these risks will accumulate across each curve used in the pricing process for a given contract.

The non-diversifiable risks from the savings-estimate and the measured-deviation are a function of their respective distributions and number of data points, and can be characterized using methods from sampling theory, as shown below:

$$\begin{aligned} & \text{standard deviation of mean-savings-estimate (standard error of mean-savings-estimate)} \\ & = (\sigma / \sqrt{n}) \end{aligned}$$

where

$$\begin{aligned} \sigma & = \text{standard deviation of savings-estimate} \\ n & = \text{number of data points (projects)} \end{aligned}$$

$$\text{standard deviation of mean-measured-deviation (standard error of mean-measured-deviation)} = (\sigma_{\Delta} / \sqrt{n_{\Delta}})$$

where

σ_{Δ} = standard deviation of measured-deviation

n_{Δ} = number of data points (projects where estimated and measured data available)

However, the mean-savings-estimate risk and mean-measured-deviation risk may diversify between each other. The extent to which they diversify between each other is often a matter of judgment, but the range can be defined. The minimum of the range is based on the assumption that the size of the savings is independent of how the estimate deviates from the true savings potential, and is given by:

$$\text{Non-diversifiable risk (min)} = \sqrt{(\sigma/\sqrt{n})^2 + (\sigma_{\Delta}/\sqrt{n_{\Delta}})^2}$$

The maximum of the range is simply the sum of the two risks and is given by:

$$\text{Non-diversifiable risk (max)} = (\sigma/\sqrt{n}) + (\sigma_{\Delta}/\sqrt{n_{\Delta}})$$

Finally, the savings curve score can be defined as follows:

$$\text{savings curve score} = \frac{1}{(\text{non diversifiable risk} / \text{mean savings})}$$

The curve score is an important criterion for selecting from a variety of curves that may be available, each with different degrees of specificity (see again figure 4). A curve generated by specifying only a few attributes will contain more data points (n will be large but σ will also be large). Conversely, a curve generated by specifying more attributes will have a lower σ , but a lower n also. The curve score can thus be used to determine the optimal tradeoff between specificity and number of data points.

Given that the non-diversifiable risk essentially has to be charged as a risk premium, it reduces the contract value. One way to manage this risk, and thereby improve contract value, is to improve the curve quality by investing in obtaining more or better data – particularly for those attributes that have the greatest predictive power.

3.3 Pricing process using curves

As noted earlier, the motivation to develop actuarial curves was to use them in the context of a scaleable pricing process. Pricing process in this context refers to the process of determining the value of potential energy savings in a portfolio of customer sites. Scalable in this context refers to the ability to dramatically increase the volume of pricing without a proportionate increase in pricing resources (people, time). While a comprehensive description of this process is beyond the scope of this paper, we describe the key aspects of this process, illustrating the use of the curves. Figure 5 contrasts the traditional pricing process with an actuarial pricing process using curves.




Figure 5. Comparison of traditional pricing process (less scaleable) to actuarial pricing (more scaleable)

Step 1 - Collect customer parameters: In an actuarial approach, instead of site audits, a standard set of data is collected from the site using some combination of email, web, and phone. One issue in this process is the veracity of customer-provided information. This risk can be hedged either contractually (by having the customer responsible for assumptions stated in the contract) or by applying appropriate risk premiums to the value of the portfolio.

Step 2 - Determine energy savings, implementation costs, time: In the purely actuarial approach, "standard" curves will be drawn from the actuarial database, based on the ECM type and parameters (location, etc.). However, a paucity of data in the actuarial database will preclude a fully actuarial approach for some ECMs. In such cases, a hybrid approach can be adopted - wherein a limited set of customer sites are audited and assessed, and curves can then developed from this data and applied to the portfolio.

Step 3 - Determine value: The actuarial process for determining the value of energy savings portfolio mimics the valuation process for energy commodity, i.e. using pricing curves as inputs to a pricing model (Figure 6). The pricing model can be standardized and much of the valuation process can even be automated, dramatically reducing pricing cycle time and errors when compared to the traditional approach with its custom pricing models.




Figure 6. Standardized ECM pricing model using ECM curves.

3.4 Use of actuarial pricing at Enron

At Enron Energy Services, the actuarial pricing approach was utilized to varying degrees in several contracts – it was most commonly used for lighting retrofits, upgrades to packaged HVAC unit replacements, and certain compressed air system measures. In at least four cases, pricing was done exclusively using an actuarial approach. In fact, plans were underway to develop standard products that could be sold with a “low-touch, high volume” sales process. For example, a product for the hospitality industry would have bundled 2 guest room ECMs and priced them exclusively from curves, using a limited amount of easily obtained site data (Figure 7). The collapse of Enron prevented the full implementation and application of this approach. Nevertheless, our experience suggests that it is particularly well suited for large portfolios of homogenous facilities such as retail outlets and hotels, where we believe this approach has the potential to dramatically reduce transaction costs and increase the scale of energy services.

ECM	Customer Parameters	ECM Curves
Guest Room Occupancy Sensors	1. # of Rooms 2. Occupancy Sensors? 3. Location 4. Occupancy Rate	1. Δ kWh/room 2. \$/room
Guest Room Lighting Upgrade	1. # Rooms 2. Incandescent Lamps? 3. Location 4. Occupancy Rate	1. Δ kW/room 2. Annual hrs/room 3. \$/room

Figure 7. Example of customer parameters and ECM curves used to actuarially price two ECMs in hotel facilities.

4 Implications for risk analysis and commodity markets

Enron’s experience with actuarial pricing of energy efficiency and its nascent attempts to commodify energy savings provides valuable lessons for the energy efficiency business, in at least three areas: risk analysis, risk management products, and tradable commodities, which we discuss below.

4.1 The need for quantitative risk analysis

From an investor perspective, arguably the most glaring deficiency in the energy efficiency business is the lack of quantitative risk analysis i.e. risk analysis models that unbundle and explicate the individual risks, characterizing each of them and their interactions with each other. Such a framework is a necessary prerequisite to the risk-return analysis that forms the basis for all investment decision-making. To investors, assessing return without analyzing associated risks is a meaningless exercise. Yet it is still not uncommon to evaluate energy efficiency projects almost exclusively in terms of simple-payback or net present value, with little if any quantitative analysis of the associated investment risk. Energy managers and investment decision-makers simply do not speak the same language (Mills et al. 2003). In the absence of quantitative risk analysis, investment decisions are disproportionately influenced by perceived risk, especially when other investments are competing for the same capital. If the risks and returns from energy efficiency projects are analyzed and evaluated similar to other investments (using well-established financial risk analysis methods) they can then be fairly compared against those other investments, and would arguably improve the chances of attracting more investment capital. For example, Rickard et al. (1998) use a common risk-return framework to compare an investment in EnergyStar homes with other non-energy investments. Risk analysis is also crucial to discounting process in evaluating energy efficiency projects (Thompson 1997).

At Enron, the pursuit of actuarial pricing in effect forced the application of financial risk analysis methods to value energy efficiency projects. Our experience suggests that this is not so much a

technical challenge as it is a professional cultural one. While energy engineers are able to make qualitative assessments of risk, they do not typically think in terms of uncertainty and probability distributions for the parameters that affect the savings and cost of a project – yet such distributions are the building blocks for quantitative risk analysis. This professional cultural barrier is not insurmountable – expert elicitation and uncertainty analysis techniques have been developed and applied in other fields whose practitioners did not previously use quantitative uncertainty analysis (Morgan and Henrion 1990). But it does require “buy-in” from these practitioners, and the necessary management processes to support it. From a technical standpoint, perhaps the most significant challenge is the paucity of data to support risk analysis. This requires a concerted public-private partnership to collect, store, and analyze project data – in effect to create industry actuarial databases, similar to what we have discussed in this paper. Other technical challenges are not as significant, given that financial risk analysis techniques are very well established, and there are several commercially available tools to support it.

4.2 The opportunity in risk management products

The development of an appropriate risk analysis framework sets the stage for the development of risk management products for energy efficiency investments. Such products are a business opportunity for entities that sell risk management services to customers who in turn benefit from the risk mitigation these products afford.

As we have shown in this paper, a key product opportunity is portfolio-based risk management. This is especially relevant given the consolidation in numerous market segments, leading to large real-estate portfolios. With the potential for portfolios of hundreds or even thousands of energy efficiency projects, portfolio-based risk management in these contexts can dramatically reduce transaction costs (e.g. site audits) as well as improve the financial viability of more risky projects (assuming those risks are diversifiable).

Actuarial databases for energy efficiency projects also support the development of energy savings insurance (ESI). ESI has considerable untapped potential to transform the market for energy efficiency by providing cost-effective risk management (Mills 2003).

There are also technological approaches to risk management. For instance, continuous commissioning can ensure that projects savings are in fact being realized, and can even increase the amount of savings realized (Claridge et al. 1994). Currently, it is very difficult to project the value of continuous commissioning for a specific site, especially without a detailed site audit. This results in a high uncertainty of the value for a specific site. However, a portfolio-based approach to continuous commissioning can reduce the overall uncertainty for the portfolio. In addition, the use of a standardized commissioning process for all sites in the portfolio can reduce the commissioning cost per site.

4.3 Energy efficiency as a commodity?

The development and maturation of risk analysis and risk management products are steps toward the evolution of a commodity market for energy efficiency. We need to distinguish here between two different commodities: the first is the direct value from energy savings (i.e. reduced energy expenditures). The second is the value from the CO₂ reductions as a result of saved energy (this commodity is often referred to as white certificates).

The primary motivation for creating a commodity market for energy savings is the premise that large-scale and efficient access to capital can be achieved through structured financial markets (Katz, et al. 1996). White certificate trading, on the other hand, is primarily a flexible means of meeting public goals for green house gas reductions (Harrington 2002). At least two countries have adopted some form of white certificate trading (International Energy Agency 2002).

While a more exhaustive treatment of this subject is beyond the scope and purpose of this paper, we would like to highlight the key implementation-related issues, based on Enron's nascent efforts to commodify energy savings.

The most significant challenge to commodifying energy savings relates to standardization of energy efficiency projects. Energy efficiency projects are remarkable for their lack of standardization, even in widely replicated projects such as office lighting retrofits. Clearly, standardization in energy efficiency projects presents a particular challenge, given the complexity and context-specific nature of most projects. However, standardization is at the core of any effort to create a commodity. In the case of energy efficiency projects, that means standardizing the way projects are priced, contracted, and implemented. This requires the development of standard project types, scope, savings measurement and verification approaches such as the International Measurement and Verification Protocol (IPMVP 2001), etc.

This presents both technical and organizational challenges. On the technical side, the challenges mainly have to do with the context specific parameters that affect the value of energy savings projects. Some are easy to characterize and treat in a standard way (e.g. weather, building area), while others are almost impossible (e.g. quality of operation and maintenance). The degree to which a particular type of energy efficiency measure can be standardized, and therefore commodified, is thus largely a function of the number and significance of parameters that can be standardized for that measure. Examples of measures that lend themselves to standardization are hotel lighting retrofits, packaged HVAC unit upgrades in retail chain facilities, and motor upgrades in constant-load industrial applications. At the other end of the spectrum are projects that are practically impossible to standardize, such as complex industrial process chiller retrofit projects.

On the organizational side, the main challenge is to develop a critical mass of stakeholders to collectively adopt the standards. This presumes that the stakeholders are interested in participating in a commodity market for energy efficiency savings. At Enron, the strategy was to simply define its own standards and use its market dominance to drive the commodification of energy efficiency, much like it did for natural gas. Without a large market player, this would require a concerted effort by the ESCO industry, possibly working through its trade associations. Commodification could also be driven by customers, financial institutions, or other stakeholders. For example, government agencies with large portfolios of similar facilities (e.g. the Postal Service) could create a demand for standardized, portfolio-based approaches to energy efficiency projects, in order to reap the benefits of commodification. Lending institutions such as the World Bank provide an impetus because commodification allows for easier aggregation of the energy savings projects, which are individually too small to merit their attention (Harris 2003). Ideally, all the stakeholders would work within the framework of a standard-setting organization, with technical and organizational support from the energy efficiency research community.

5 Acknowledgements

The authors would like to acknowledge their former colleagues at Enron Energy Services who collaborated on the work described in this paper, particularly Nancy Simmons, David Truncale, and Sam Zimmerman who worked on implementing the actuarial database. The authors also thank Jeff Harris at Lawrence Berkeley Laboratory for his detailed review and insight on policy implications of this work.

6 References

- ASHRAE (2002). *Guideline 14: Measurement of Energy and Demand Savings*. American Society of Heating Refrigeration and Air-conditioning Engineers, Inc., 1791 Tullie Circle, NE, Atlanta, GA 30329-2305.
- Chamberlin, J., Herman, P. (1995) "The Energy Efficiency Challenge: Save the Baby, Throw out the Bathwater," *The Electricity Journal*, December 1995. 38-47.
- Claridge, D.E., J. Haberl, M. Liu, J. Houcek, A. Athar. (1994) "Can you achieve 150 percent of predicted retrofit savings? Is it time for recommissioning?" *Proceedings of the ACEEE 1994 Summer Study on Energy Efficiency in Buildings*, Vol. 5. American Council for an Energy Efficient Economy. Washington D.C.
- Goldman, C, P. Juergens, M. Fowlie, J. Osborn, K. Kawamoto, T. Singer. (2000) "Historical Performance of the U.S. ESCO Industry: Results from the NAESCO Project Database", *Proceedings of the ACEEE 2000 Summer Study on Energy Efficiency in Buildings*, 5: 89-102. American Council for an Energy Efficient Economy. Washington D.C.
- Harrington, P. (2002) Energy Efficiency Certificate Trading, Presentation at IEA/DSM workshop, Milano, April 17, 2002.
- Harris, J. (2003) Personal Communication. Lawrence Berkeley National Laboratory. August 22, 2003.
- International Energy Agency (2002), *DSM Spotlight*, May 2002. <http://dsm.iea.org>.
- International Performance Measurement & Verification Protocol (2001). "Concepts and Options for Determining Energy and Water Savings", US Department of Energy, DOE/GO-102002-1554.
- Kats, G. H., Rosenfeld, A. H., McIntosh, T. A., McGaraghan, S. A. (1996). "Energy Efficiency as a Commodity: The Emergence of an Efficiency Secondary Market for Savings in Commercial Buildings." *Proceedings of the ACEEE 1996 Summer Study on Energy Efficiency in Buildings*. American Council for an Energy Efficient Economy. Washington D.C.
- Mills, E., Kromer, S., Weiss, G. (2003). "From Volatility to Value: Analyzing and Managing Financial and Performance Risk in Energy Savings Projects", *ECEEE Summer Study*, 2003
- Mills, E. (2003). "Risk Transfer via Energy Savings Insurance". *Energy Policy* 31 (2003) 273-281
- Morgan, G., M. Henrion, (1990). *Uncertainty*. Cambridge University Press.
- Office of Industrial Technologies (2003). Industrial Assessment Centers Database [IAC 2003]. U.S. Department of Energy. <http://www.oit.doe.gov/iac/tools.shtml>.
- Reddy, T.A., Haberl, J.S., Elleson, J.S. (1999). "Engineering Uncertainty Analysis in the Evaluation of Energy and Cost Savings of Cooling System Alternatives Based on Field-Monitored Data" *ASHARE Transactions V. 105*, pt2.
- Rickard, S., B. Hardy, B. Von Neida, P. Mihlmester. (1998). "Investment Risk in Whole Building Energy Efficiency Upgrade Projects". *Proceedings of the ACEEE 1998 Summer Study on Energy Efficiency in Buildings*, 4:307-318. American Council for an Energy Efficient Economy. Washington D.C.
- Sioshansi, F. (2002) "The emergence of trading and risk management in liberalized electricity markets," *Energy Policy* 30 (2002) 449-459.
- Thompson, P.B. (1997). "Evaluating energy efficiency investments: accounting for risk in the discounting process," *Energy Policy*, Vol. 25, No.12, pp. 989-996, 1997.