POWER

PWP-028

THE ECONOMICS OF CONSERVED-ENERGY "SUPPLY" CURVES

Steven Stoft

April 1995

This paper is part of the working papers series of the Program on Workable Energy Regulation (POWER). POWER is a program of the University of California Energy Institute, a multicampus research unit of the University of California, located on the Berkeley campus.

University of California Energy Institute
2539 Channing Way
Berkeley, California 94720-5180
www.ucei.berkeley.edu/ucei

PWP-028

THE ECONOMICS OF CONSERVED-ENERGY “SUPPLY” CURVES

Steven Stoft

April 1995

Table of Contents

The Economics of Conserved-Energy "Supply" Curves

The Economics of Conserved-Energy "Supply" Curves *

Steven Stoft **

10 January 1996

This paper develops the theoretical underpinnings of conservation "supply" curves (CSCs), and in doing so uncovers several problems with current procedures for their construction. The CSC is shown to be derivable from a production isoquant, and not to be a true supply curve. The traditional algorithm for constructing a CSC from discrete measures is shown to be suboptimal, contrary to prior claims. Omitting conservation measures from consideration can lead to systematic, excessive conservation. The CSC concept is extended from constant-service to constant-utility measures, and an improved approximation is suggested for the cost of conserved energy (CCE) of measures that cause rebound. The appendix provides a formula for CCE that is simple yet more general than the one currently in use, but shows that even with this generalization, CSCs cannot be constructed for a world with fluctuating energy prices.

INTRODUCTION

Conservation supply curves have now been a primary analytic and advocacy tool of the energy efficiency community for over a decade. Parties engaged in their construction and use have included, the National Academy of Sciences,[1] the U.S. Department of Energy,[2] the New England Energy Policy Council (1987), and the Electric Power Research Institute (1990), and their use has been strongly advocated by such prominent figures as Arthur Rosenfeld and Amory Lovins. In spite of such prominence and in spite of their claim to integrate the engineering

---

* I would like to thank Carl Blumstein for insisting that a CSC is not a normal supply curve, and asking how a CSC should be defined. I would also like to thank him for many useful hours of consultation on these topics, and for providing much of the background used in the introduction. I am also indebted to Chuck Goldman, Alan Meier, Jon Koomey, Severin Borenstein, Bart McGuire, and Haru Connolly for many useful comments and questions.

** University of California Energy Institute, 2539 Channing Way, Berkeley, CA 94720; and Lawrence Berkeley Laboratory, University of California, Berkeley, CA 94720, USA.

1. "Technological costing" was one approach used by the National Academy of Sciences' Panel on Policy Implications of Greenhouse Warming in its evaluation of options for mitigating greenhouse warming. This approach relied on the construction of a national conservation supply curve. Most of the results on mitigation in the panel's report are based on this approach. See Panel on Policy Implications of Greenhouse Warming (1992), and Rubin (1992).

2. A national conservation supply curve for residential electricity was constructed for DOE by Lawrence Berkeley Laboratory. See Koomey, et al. (1991).

approach to conservation with the economic approach, the underlying economic theory of these curves has never been presented. This paper begins to remedy that oversight, by showing how a conservation supply curve (CSC) is derived from a production function, and by presenting an optimal algorithm for constructing a CSC from discrete conservation measures. While providing a sound theoretical basis for construction and use of CSCs, these developments also point to past errors.

Conservation supply curves were introduced and popularized by Rosenfeld and his colleagues at LBL (see 1983 Meier, Wright, and Rosenfeld, 1983).[3] Subsequent work by Rosenfeld and a number of other investigators is summarized in Rosenfeld et al. (1993). The most widely publicized CSCs have been those describing conservation potential at the national level, such as the one used by the National Academy. On a smaller geographical scale, CSCs have been constructed for particular utility service territories as exemplified by the work of the Bonneville Power Authority, and Northwest Power Planning Council.[4] Often they have been constructed for a market segment, such as residential and commercial electricity use.[5] They have been used occasionally for ranking retrofit strategies for individual buildings, and even in the design of a single building.[6]

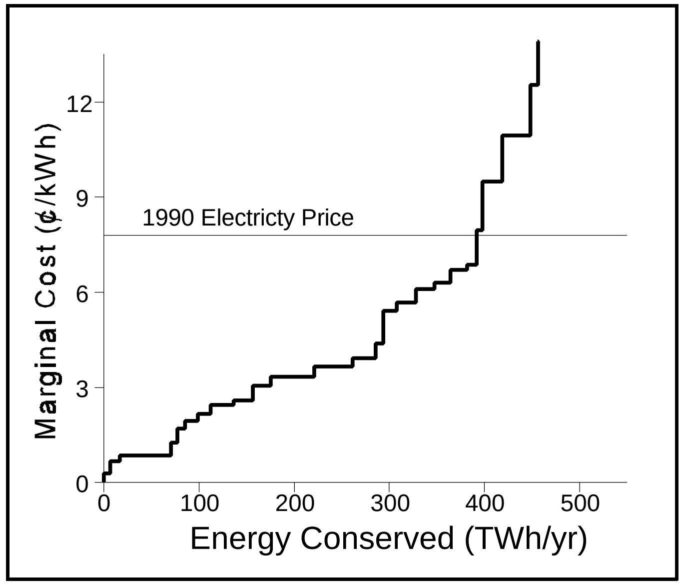

Figure 1. A Conservation Supply Curve for U.S. Residential Electricity

Although constructing a CSC for a region differs markedly from the process of CSC construction for a building, this difference lies entirely in defining and evaluating the measures. Once this has been done, and the measures have been described by their energy savings and their costs, the process of using these data to construct a CSC is the same regardless of the type of CSC. This paper focuses entirely on the problems that are common to all types of CSC, and

---

3. Rosenfeld reports that work on conservation supply curves was initially stimulated by a suggestion from Roger Sant *ca*. 1978. Work predating the publication of Meier, et. al. (1983) includes Wright's Masters Thesis, Meier's Ph.D. Thesis, and a report sponsored by the California Policy Seminar.

4. See Usibelli (1983) and Northwest Power Planning Council (1991).

5. Such CSCs have been constructed for the Electric Power Research Institute (1989), the California Energy Commission (Goldstein et al., 1990), and the New England Energy Policy Council (1987) among others.

6. This was done in the ACTT project at Pacific Gas & Electic under the direction of Grant Brohard.

especially on producing the CSC itself from the data points that describe packages conservation measures.

A CSC plots the marginal cost of conserved energy against the total amount of energy conserved, as is shown in **Figure 1**. It is based on a CSC for U.S. residential electricity usage presented by Rosenfeld et al. (1993), but it represents only stylistically the more than 200 individual measures plotted on the actual CSC. The horizontal line representing the price of electric energy, divides the measures on the CSC between those below, which provide conservation at a positive net present value (including implementation, maintenance and energy savings), and those above, which do not. This division represents a primary use of CSCs, and the ability of a CSC to divide conservation measures in this way is a defining feature.[7]

This particular conservation supply curve is quite well behaved, being both non-negative and monotonically increasing. Although we will see that these are theoretical requirements, not all published CSCs conform as required. For instance, the Rocky Mountain Institute (Lovins et al., 1986) has published one that begins with lighting measures having a marginal cost of less than −1¢/kWh, and has more recently identified a "tunneling effect" which is claimed to produce non-monotonic CSCs.

BASICS OF CONSERVATION SUPPLY CURVES

A traditional CSC is based on a sequence of packages of conservation measures, each of which provides additional conservation at a higher marginal cost than the previous packages. Typically the first package consists simply of the single most economic conservation measure, and each subsequent package adds a slightly less efficient measure. Sometimes measures conflict. Triple-pane windows and double-pane windows are mutually exclusive, so the double-pane measure will be dropped when the triple-pane measure is added in the sequence of conservation packages.

The crucial step in the construction of a CSC is the calculation of the marginal cost of conserved energy (*CCE*), which is computed by dividing the total cost of conservation (*TCC*) by the total energy savings (ΔE). The difficulty with the concept of the cost of conserved energy is in knowing to what it applies. Clearly, to compute ΔE, we must consider two production technologies, one before and one after the conservation measure. But, conservation measures are usually defined in such a way that one does not know either the starting or ending technology, but only the change in technology. For example, a measure might specify an increase in ceiling insulation from four inches to eight inches, without specifying the efficiency of the building's furnace.

---

7. "All measures on the conservation supply curve below the current price of electricity provide energy saving at a net negative cost." Rosenfeld et al. (1993). For this line to divide true net present values, the definitions of *TCC* and ΔE given in the appendix must be used.

Because of measure interactions, initial conditions generally affect a measure's ΔE. In spite of this, *CCE*s are typically said to be calculated for specific measures, as if the result were independent of the initial technology and depended only on the measure itself.

Generally, when the *CCE* of a particular measure is referred to it means one of two things, and in order to keep these meanings distinct, this paper will use two different symbols.

1. CCE_0 applies to a measure taken on top of some standard base case.[8]
2. *CCE* applies to a measure taken after the next most efficient conservation package.

The two definitions are actually very different in nature. The second definition is a definition of marginal cost in a setting where changes are small but discrete. The first definition is non-marginal in nature; with an *inefficient* base case its results can be very different than those of the marginal definition, as will be seen later in an example.

We will be concerned primarily with marginal *CCE* because that is the one mainly used when constructing a CSC. This is made clear by the CSC construction algorithm presented by Koomey et al. (1991). They state that after each measure is implemented "the new energy use is used to recalculate the CCE's of the remaining measures." Thus the CCEs graphed by a CSC do not generally represent changes from the base case (CCE_0), but instead represent marginal CCE.[9]

For now we will avoid the problems of computing CCE, not by assuming away interactions, but by assuming that we know the optimal sequence of conservation packages.[10] To construct a CSC, assign each measure its *CCE* and ΔE, based on the optimal sequence. Then choose the measure with the lowest *CCE,* and plot *CCE* against ΔE as the first point on the CSC. Then find the measure with the next lowest *CCE* and plot its *CCE* against the sum of its ΔE and all previous ΔE's. In this way the CSC is constructed point by point.

The most traditional supply curve calculations, such as those found in Meier, Wright and Rosenfeld (1983), assume that all conservation measures last for the same number of years and produce a *constant* flow of conserved energy during this time. This simplifying assumption makes it possible to use a specialized formula for *CCE* and easy to measure conserved energy as a flow of kWh/year. However, when a broader range of measures is considered, which may include differing lifetimes and more complex flows of conserved energy, the traditional *CCE* formula becomes unworkable. Fortunately there is a more general method of handling these complexities, that is as simple as the traditional method. The general method defines *TCC* as the present value of conservation costs, and ΔE as the present value of conserved energy. As is

---

8. This definition is commonly used in two distinct circumstances. In a study of a set of conservation measures that does not construct a CSC, such as the lighting study by Piette et al. (1989), this definition is often used exclusively. The definition is also used (though it should not be) in CSC construction for those measures which are judged not to interact with previous measures. These will later be defined as "first measures."

9. This is also recognized by Rubin et al. (1992), who describe a CSC (their Fig.1) as a "marginal cost curve for building sector electricity use."

10. Later it will be seen that such a simple procedure also requires that the base case be technically efficient.

pointed out by Meier (1982) a crucial property of the *CCE* is its lack of dependence on energy prices. Unfortunately, as the appendix demonstrates, in a world of fluctuating energy prices, no formula for *CCE* can both retain this property and make possible the construction of a working CSC.

Since the central concerns of this paper can be fully explained without reference to any particular method of calculating *TCC*, *ΔE*, and *CCE*, from here on we will assume that a suitable method has been adopted.

Before returning to the questions raised by the analysis of discrete conservation measures and their interactions, we explore a more classical economic approach to the problem, based on continuous productions functions. We do this by adopting the view that a CSC is simply a curve showing the cost of conserving one input to a normal production process. In this view, a CSC simply answers the well defined question, how much more does it cost to produce a fixed output level as a single input is reduced?

FROM PRODUCTION FUNCTIONS TO CONSERVATION CURVES

In order to discuss the cost of conserving something, one must define the production process for which it is an input. For example, one cannot discuss the cost of conserving steak without first defining steak to be an input to the production of the consumer's utility. Otherwise there is nothing to be held constant when steak is conserved. With nothing to hold constant, conserving one steak is always found to save the cost of the steak. Once we define the production process, we can ask an interesting question: how does the cost of producing a fixed output vary with the level of some particular input, *E*?

Fortunately, energy is used in the production of intermediate goods, such as light and heat; this allows us to work with explicit production functions. However we will see shortly that it is often useful to consider energy as an input to the production of utility, because this allows for a much wider range of substitution and therefore a greater potential for conservation.

Aggregating Non-Energy Inputs

Generally an energy service is produced by using several inputs besides energy and so must be described by a production function such as $Q(E, x_1, ..., x_n)$, where $x_1, ..., x_n$ are the non-energy inputs.[11] Because we wish to find the cost of conserving *E*, we must summarize the non-energy inputs by a single cost, which should be the minimum possible cost for a given level of output. This means we are looking for a production function of the form $Q(E, C)$, which gives

11. Many of the ideas presented here concerning the relationship of CSC to productions functions are also presented by Blumstein and Stoft (1995) in the context of a discussion of Huntington (1994).

the maximum possible output from input, *E*, when cost, *C*, is spent optimally on all other inputs. This function is defined by:

$$Q(E, C) = \max_{\{x\}} Q(E, x_1, ..., x_n) \quad \text{s.t.} \quad \sum_{i=1}^{n} p_i \cdot x_i = C \qquad \textbf{(1)}$$

Since we will be holding output fixed, plotting an isoquant proves illuminating. In this case we wish to plot the isoquant that corresponds to the level of output before the conservation measures under consideration have been implemented. Such an isoquant is shown in **Figure 2**.

Characteristically this isoquant is downward sloping except for extremely high values of energy input. At these values, more cost may be incurred trying to dispose of waste energy than is saved by wasting the energy. In other words this occurs only when the standard economic hypothesis of "free disposal" is violated. Conditions under which this occurs are discussed at the end of the appendix. Put in terms of "implicit discount rates," this part of the curve represents a discount rate beyond infinity. Although design errors may lead to such situations, it seems unlikely to find the base case in this region of the isoquant. Generally we will assume that it is not.

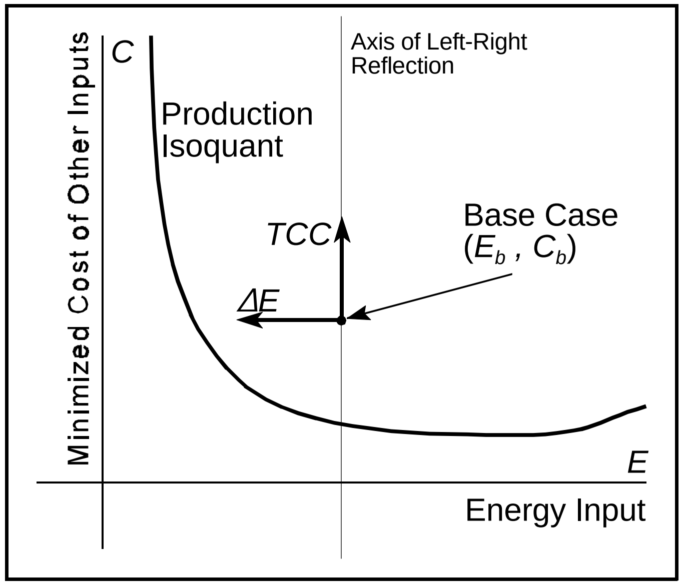

Figure 2. Isoquant of *Q*(*E*, *C*)

From Production Isoquant to CSC

The base case, shown in **Figure 2**, must be chosen to provide a reference against which to measure conservation. The base case is defined by its coordinates (E_b, C_b), and by its output level *Q*. This is the same *Q* for which we draw the isoquant, so if the base case is technically efficient, it will lie on the isoquant. If not, it will lie above the isoquant as shown in **Figure 2**. This can happen because of a market failure, because the base case represents technology that was not available when the isoquant was constructed, or because the CSC is being used for optimizing a new design and the base case is just an arbitrarily chosen reference point.[12]

12. Although most supply curves describe what are supposedly, and may be, rather extreme market failures, this paper takes a completely agnostic view on the existence of such failures. The construction of CSCs neither depends on their existence, nor is particularly useful in proving they do exist. When a CSC intersects the vertical axis well below the price of energy, one can believe either that figures it is based on were correctly calculated, or that costs were omitted and savings overstated. But, nothing in the procedure for drawing the curve will shed any light on this question.

The base-case point serves as the origin when graphing the total-cost-of-conservation curve (TCC curve), which is defined by the left-right reflection of the isoquant in the vertical line through that point, as shown in **Figure 3**. Thus the vertical axis measures the amount of energy conserved relative to the base case, and the horizontal axis measures the cost of conserving that energy. Analytically, the TCC curve is defined by specifying that if (E,C) is a point on the isoquant, then (ΔE, TCC) is a point on the TCC curve, where $\Delta E = E_b - E$, and $TCC = C - C_b$. Note that the TCC curve is monotonically increasing, simply as a result of its being a left-right reflection of the production isoquant, which is monotonically decreasing (except as noted at the extreme right).

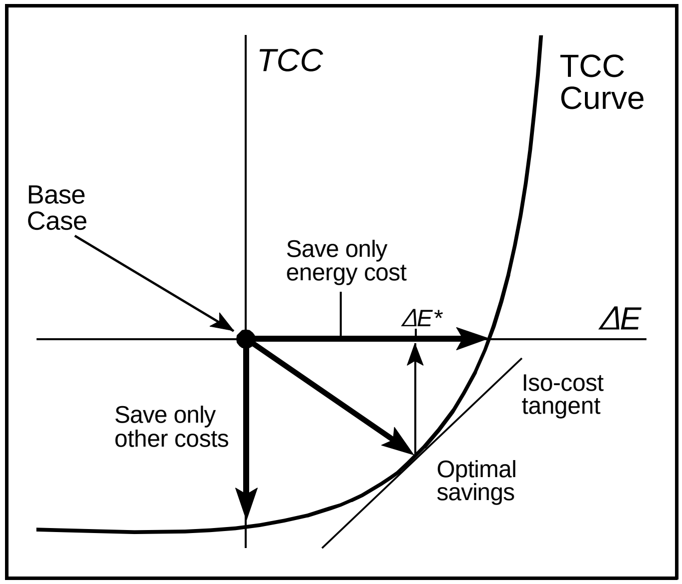

Figure 3. Total-Cost-of-Conservation Curve

The TCC curve displays the *total* cost of conserving energy, but, for efficiency to prevail, it is the *marginal* cost of conserving energy that must equal the marginal cost of producing energy. The CSC, shown in **Figure 4**, plots this marginal cost and consequently is constructed simply by taking the derivative of the TCC curve; $CCE(\Delta E) = TCC'(\Delta E)$. Because a CSC is the derivative of the monotonically increasing TCC curve, it will be positive everywhere except for extremely high levels of energy use, which generally only occur well to the left of the origin (base case). (This is the same exception as was discussed for the production isoquant.) Since the TCC curve is defined for both positive and negative values of conservation, so is the CSC. Thus a normal CSC will cross the *CCE*-axis at a positive value and continue to the left of this axis as shown in **Figure 4**.

Note that because the base case can be above the isoquant, a TCC curve may be negative when $\Delta E = 0$. This is because, an inefficient base case makes it possible to produce the same output (of energy services) using the base level of energy but a less costly combination of other inputs. However this does not indicate that the CSC will be negative.

If the base case is efficient then it must lie on the production isoquant. It also must be located at the point on the

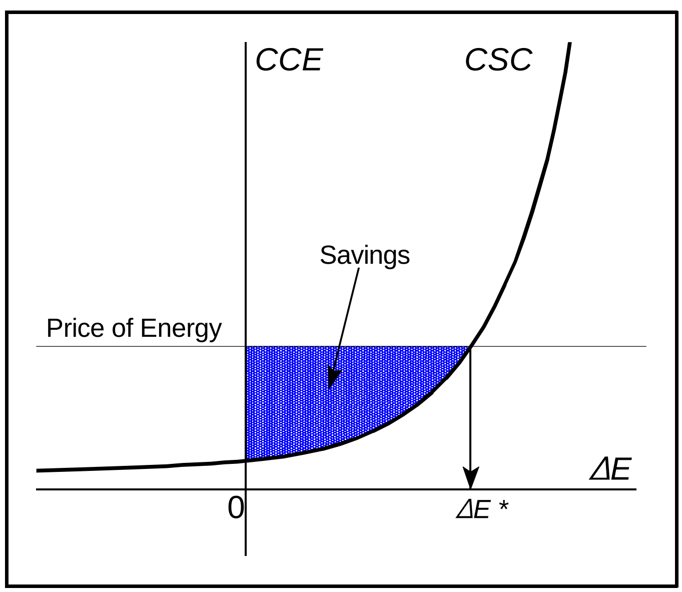

Figure 4. A Conservation “Supply” Curve

production isoquant where motion along the isoquant decreases the cost of energy at the same rate that it increases the cost of other inputs, otherwise money can be saved by moving one way or the other. At this point, the slope of the production isoquant is minus the price of energy. So, with an efficient base case, the height of the CSC at the origin will be exactly the marginal cost of energy.

Interpreting the TCC Curve and the CSC

Although the TCC curve and the CSC describe the same set of conservation possibilities, they each have their own advantages. The TCC actually contains more information than the CSC, since the CSC is derived from it and since it is not possible to reverse this operation and derive the TCC curve from a CSC. This is because information about the TCC curve's vertical position is not deducible from its slope. Eliminating information about the TCC's vertical position can actually be seen as an advantage because it prevents users of a CSC from attributing to energy conservation the savings from eliminating non-energy inefficiency. To illustrate this point we consider the technically inefficient example diagramed in **Figure 2** through **Figure 4**.

If the base case is technically inefficient, as is illustrated in **Figure 3**, then moving onto the technically efficient isoquant will save money even when energy use is held constant. This possibility is shown by the downward vertical arrow. This savings should not be attributed to energy conservation because it involves no change in energy use. Similarly it is possible to save energy without any change in the cost of other inputs. This raises the questions of how much energy should be conserved, and how much savings should be attributed to energy conservation.

The answer is that total cost, including the cost of energy and all other inputs should be minimized, and that this can occur anywhere on the TCC curve depending on the price of energy. Thus it is possible that other costs should be increased in order to save even more energy, or that energy use should actually be increased above its base case level. In the example of **Figure 3** the iso-cost line is tangent to the TCC curve at a point south-east of the base case. Since total cost decreases to the south-east, the minimum cost point is defined by the equality of the slope of the TCC curve and that of the iso-cost tangent. This slope is just the price of energy. In order to help identify this point of equality, the CSC plots the slope of the TCC curve. The level of optimal conservation is then found by superimposing a horizontal line at the price of energy and noting its intersection with the CSC, as shown in **Figure 4**.

The CSC graph now provides a straightforward answer to the question (which appeared confusing in the context of the TCC curve) of how much money is saved by conservation. When 1 kWh of energy is conserved at the marginal cost of conserved energy, then the savings is clearly $P - CCE$, where P is the price of energy. Thus the maximum total savings from conservation is simply the shaded area of **Figure 4** between the energy price line and the CSC, and between zero and ΔE^*, where ΔE^* is the optimal level of energy conservation (see Rosenfeld et al., 1993).

Is it a True "Supply" Curve?

Now that a CSC has been defined by its relationship to a production function, we can answer the question: Is a CSC a supply curve in the normal sense of that term? Since its relation to a production function is quite different than that of a normal supply curve, the answer is no. (In fact a CSC is actually an unusually oriented conditional factor-demand function.[13]) But for a deeper understanding of how it is different, and the importance of this difference, we explore the possibility of defining a true supply curve of conserved energy.

If we interpret the marginal cost, *CCE*, as a price, then a CSC resembles a true supply curve since it relates the supply of conserved energy to a price for that energy. However, when we try to push the analogy to completion, we must look for a supplier and ask how that supplier would behave if offered various prices for conserving energy. This thought experiment will help us distinguish between a CSC and a standard supply curve.

For simplicity let us assume we have correctly computed a CSC for a single neoclassical consumer who has already optimized energy use at the prevailing price of 10¢/kWh. Now we approach this consumer with an offer to pay 10¢ for every kWh of conservation supplied. How much will she conserve? The answer is: quite a lot, because having previously optimized energy consumption, she is at the margin exactly indifferent between consuming and not consuming her last kWh. Being indifferent means that 10¢/kWh is much more than enough to induce conservation.

This indicates a "standard supply curve" that is far to the right of the origin at the price of 10¢/kWh; in fact it is to the right of the origin at any price above zero. Does this mean society should buy conservation at any price up to 10¢/kWh? No, as has been argued at length by Ruff (1988), this ignores the consumer's reduction in payment to the utility;[14] optimizing consumers have not made a mistake, and society should not pay even 1¢/kWh in such a situation. Fortunately, a CSC does not tempt us into such mistakes, for as mentioned earlier, it will cross the vertical axis at 10¢/kWh. So a CSC is not a true supply curve, and this is all for the better.

The confusion between CSCs and true supply curves, may have helped legitimize the view that the supply of conservation should be treated as a normal supply process. This view has lead to the introduction of the concept of the "negawatt," which is said to measure conserved energy (Lovins, 1985). Arguably, this misconception has created confusion in the "level playing field" debate over the proper rules for all-source bidding, a process by which utilities buy watts and negawatts in a competitive auction.

---

13. If $E(P, p, Q_b)$ is the conditional factor demand for energy, then $\Delta E = E_b - E(P; p, Q_b)$ with p and Q_b considered parameters, defines the CSC. P is the price of energy, p is the price vector for other inputs (x), and Q_b is the output level of the base case. See Varian (1978, p. 31).

14. When the consumer gives up 1kWh, she is paid for conservation, say 5¢, and pays the utility 10¢ less for a total gain of 15¢. In return she "supplies" 1kWh of conserved energy which is "used" by the utility to reduce generation. Thus she is effectively paid retail price of energy plus the payment for conservation.

CONSTRUCTING A CSC FROM DISCRETE MEASURES

Generally conservation supply curves are not constructed from analytic production functions of the type used in the above analysis, but from data on discrete conservation measures. This section considers how to construct a CSC from a list of measures and their attributes, paying special attention to the relationship between this construction procedure and the theory presented above. In keeping with the above analysis and also in keeping with standard practice, we will consider only combinations of conservation measures that leave the consumer with a level of utility that is unchanged from the base case. Thus we are only concerned with measures that maintain a constant energy-service level.

Unfortunately there is one complication that must be faced immediately, and that is interaction between measures. If two measures *A* and *B* are both implemented, the combined measure *AB* will not save as much energy as the sum of the two individual measures' savings. Generally the cost of *AB* will be the sum of the cost of *A* and *B* conducted separately, but even this is not necessary. A typical type of interaction is that between insulation and increasing the energy efficiency of a furnace: the more insulation, the less energy savings from an improvement in the furnace. In **Figure 5**, showing interactions between four common areas for conservation, an arrow from *A* to *B* indicates that an efficiency improvement in *A* changes the ΔE of doing *B* after *A*.

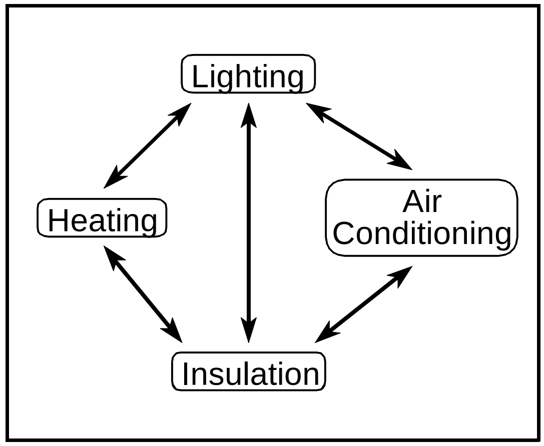

Figure 5. Conservation Interactions

As a consequence of these interactions it is necessary to collect data not just on individual measures but on nearly all combinations of measures. The only combinations of measures that can, and in fact must, be omitted are those that cause a change in utility from the base case. For the sake of brevity we will refer to individual measures and combinations of measures simply as packages. This motivates a pair of definitions.

- A set of measures, M, is said to be **well defined** relative to some base case if and only if it is possible to compute *TCC* and ΔE for every subset of M, where *TCC* and ΔE account for total cost of conservation and total energy savings respectively.
- A **package** is any subset of M.

From now on, whenever we speak of a measure or a package, we will mean one that is part of a *well-defined* set of measures. A package can be represented by the pair (ΔE, *TCC*), and the appendix describes exactly how these values should be computed from the relevant cost and energy streams.

A Suboptimal Algorithm

The method of constructing a CSC from conservation packages has generally been described as follows. Find the single most economic (lowest CCE) individual measure, call it *A*, and plot it first (ΔE_A, CCE_A). Then consider all packages that include *A* and any other single measure, and find the most economical of these packages. Call this package *AB* and plot it second, (ΔE_{AB}, CCE_{AB}). This process is continued until all of the measures have been included in the final package, and it produces what we will call a ***sequential ordering*** of the individual measures. Note that this concept of sequential ordering rules out some possibilities, for instance: *A* followed by *B* followed by *AB* is ruled out.

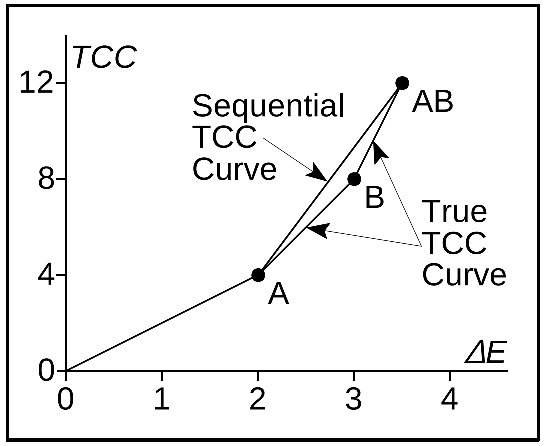

Figure 6. Sequential and Correct TCC Curves

In what has probably been the only attempt to rigorously specify the optimal algorithm for constructing a CSC, Meier (1982, p. 18-19) specified the algorithm just described and claimed that it produces the "optimal investment schedule." A decade later this algorithm was still in use and was still thought to be optimal (Koomey et al., 1991). Unfortunately, the optimal ordering of conservation packages is not always sequential, so Meier's algorithm cannot always be optimal. To demonstrate this consider the following two measures shown in **Figure 6**: $A = (2, 4)$ and $B = (3, 8)$. The first is an improved air conditioner, the second is improved insulation that nearly eliminates the need for air conditioning. Because of this interaction, the combination measure *AB* equals (3.5, 12), indicating that the costs are additive, but the energy savings are subadditive, $3.5 < 2+3$.

As can be seen from **Figure 6** the sequential TCC curve is not dramatically different than the true TCC curve. However the small difference in the position of these curves causes a much more significant difference in their slopes, so the related CSCs differ significantly. In **Figure 7** each horizontal line is labeled with the change for which its *CCE* is computed. Two different *CCE*s are associated with *AB*, one when it follows *A* and the other when it follows *B*.

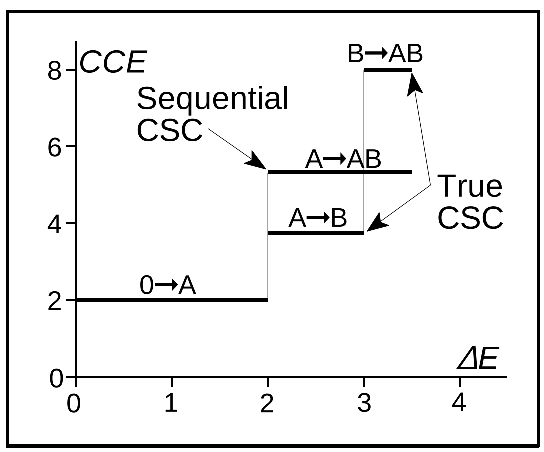

Figure 7. Sequential and Correct CSCs

Notice that according to the sequential CSC, the *AB* package becomes economical below 6¢/kWh, while the true CSC shows that it should not be used unless the price of energy exceeds 8¢/kWh. Conversely, if the price of energy is between 4 and 5¢/kWh, the sequential CSC indicates that only measure A should be taken, while in fact it is measure B that is appropriate and saves more energy.

This example demonstrates that the standard sequential algorithm is suboptimal and that the optimal ordering of conservation packages is generally not a "sequential order" of individual measures. Additionally it can be shown that Meier's algorithm can generate CSCs that have decreasing regions. In the following discussion of an optimal algorithm, we will show that correctly constructed CSCs are monotonically increasing, as we found in the previous section.

An Optimal Algorithm

Although the previous example made use of the correct CSC construction algorithm, we have not yet specified that algorithm precisely nor have we argued for its validity. The correct algorithm must satisfy two conditions. It must select a set of conservation packages, {CP}, such that

1. All packages in {CP} are justifiable at some energy price level.
2. No omitted package is preferred to all packages in {CP} at any single energy price.

To respect criterion 2 we begin by graphing all conservation packages that satisfy our criteria of no change in utility from the base case.[15] The next step is to construct the convex lower bound of this set. A package is a vertex of this bound if there exists a tangent line through it such that all other points lie above that line. The bound itself consists of these vertexes and the straight line segments that connect them sequentially from left to right. This situation is shown in **Figure 8**.

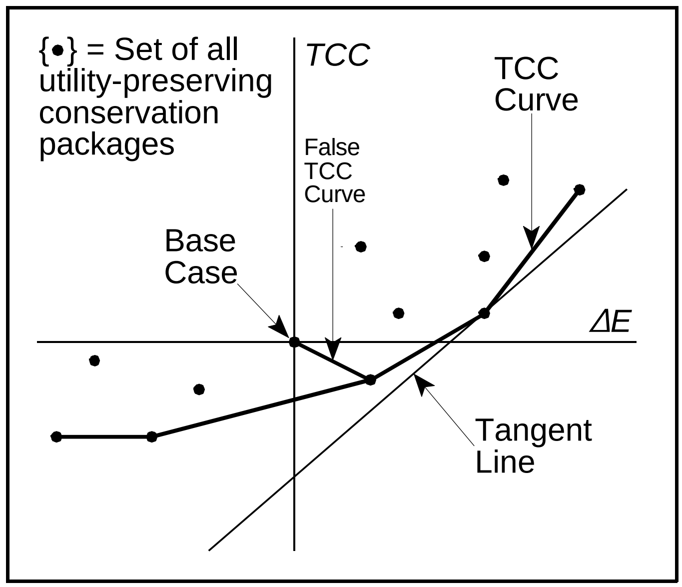

Figure 8. A Discrete TCC Curve

Note that each point on the lower convex bound is the most economic choice at the price corresponding to its tangent line; thus condition 1 is satisfied. Note also that if a package is preferred to all others at some price, then the iso-price line passing through that package's point will be a tangent line, and thus the package will be part of the lower convex bound. Thus condition 2 is satisfied. Thus the packages that are part of the lower convex bound define the TCC Curve.

15. If we had not changed axises, but were still working with *E* and *C*, this set of packages would correspond exactly to the "input requirement set" used in constructing the isoquant of a production function (except for the fact that we have replaced non-energy inputs with their minimum cost, C). See Varian (1978) page 4. Our construction of the TCC curve corresponds to the standard definition of an isoquant, except that only the nodes would be included if the other points on the TCC could not be obtained through mixing productions strategies. Constructing a *TCC* is exactly the same as constructing a production isoquant as a production possibility frontier.

Having found the TCC curve, constructing the CSC is merely a matter of plotting the slopes of the TCC curve's straight line segments. Notice that because the TCC curve is a lower convex boundary, its slope is monotonically increasing, which means that the CSC is also monotonically increasing.

Figure 8 demonstrates another problem with the standard sequential algorithm for constructing CSCs. The algorithm specifies that the first measure's CCE should be assessed against the base case. But this is the non-marginal definition of CCE, which we have termed CCE_0. If we were only concerned with positive-conservation measure *and* if the base case were technically efficient, there would be no problem. This is because, for a technically efficient base case, the TCC curve passes through the origin. But in our example, the standard algorithm produces a negative-CCE first measure, as is indicated by the segment of TCC labeled "False TCC Curve." As was mentioned earlier, the Rocky Mountain Institute (1991) has published CSCs that begin at negative levels of CCE, and this may be the explanation.[16]

Problems Caused by the Omission of Measures

If a CSC is used correctly to determine a conservation investment strategy, that strategy will lead to improvement, even if some measures are omitted. But omitting measures obviously can lead to suboptimal gains. Interestingly there are at least two types of omissions that may occur systematically and that may lead to excessive conservation. The first type is the omission of low-conservation or negative-conservation measures, and the second type is the omission of partial-program measures. We will consider each in turn.

To demonstrate the possible effects of omitting low- and negative-conservation measures, we analyze an example demonstrating how such omissions could lead to serious underestimation of the CCE for "first measures." A "first measure" is defined as one whose *CCE*, as reported in the supply curve, is the same as if measured from the base case ($CCE = CCE_0$). By this definition, there may be quite a few "first measures" in any given CSC.

Consider a first measure, consisting of fluorescent lighting, called *A*, with a slightly negative CCE_0, say −1.25¢/kWh. The negative non-energy cost of such a measure might be due to the dramatically reduced maintenance cost associated with longer bulb life. But perhaps there is a measure that saves on maintenance and is even cheaper to install, but uses more energy than the status-quo base case. This might be a long-life, low-efficiency incandescent lighting

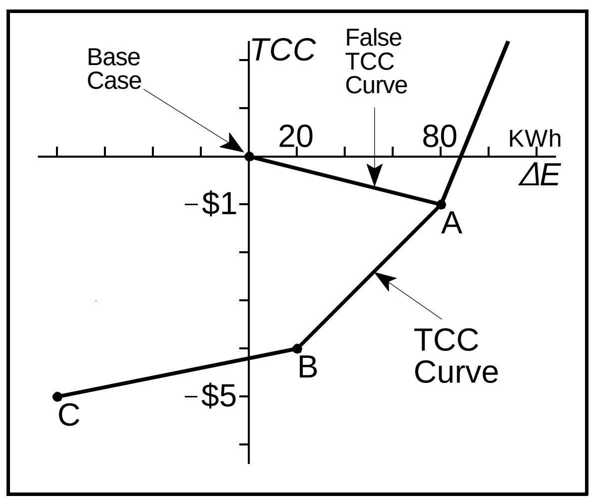

Figure 9. The Role of Anti-Conservation

16. See Joskow and Marron (1992, p. 49) for a similar explanation.

system.[17] We know such a measure would be omitted from the CSC because negative-conservation measures are systematically excluded. Call this measure *C*. Since negative-conservation measures are being systematically excluded, there may also be a tendency to omit low-conservation measures, especially when high-conservation measures are readily available. Let us suppose that there is such a measure and call it *B*. It might look a lot like *A*, but be a cheaper and less energy efficient fluorescent. These measures are displayed in **Figure 9**.

The first point to notice concerning **Figure 9** is that while the CCE_0 of measure *A* is −1.25¢/kWh, it's marginal *CCE* is 5¢/kWh. Thus, as previously noted, the distinction between these two definitions has very real consequences. Second, note that according to the false TCC curve, measure *A* would be worthwhile at any price of energy, but that the true TCC curve indicates it is worthwhile only if the price of energy is below 5¢/kWh. Thus omitting low- or negative-conservation measures can lead to over conservation. Next notice that the CCE_0 of measure B is −20¢/kWh. If this were mistaken for the measure's *CCE*, we would conclude that the measure is an extraordinarily successful energy conservation measure. But the truth is that, while this measure is very cost effective, it savings should not be attributed entirely or even primarily to energy savings. This is primarily a maintenance saving measure, and its true *CCE* is a very reasonable +1¢/kWh. Thus it saves about $4 in maintenance, and, at an energy price of 8¢/kWh, it saves $1.40 in energy costs (not (8+20)¢/kWh·20kWh = $5.60 in energy costs).

Note that if the price of energy is 8¢/kWh, then measure *A* would be adopted under either construction of the supply curve. The only difference would be that under the erroneous construction, it would be credited with considerably more energy savings. In this case all of the maintenance-cost savings could be attributed to energy-cost savings. This is not such a consequential mistake, but clearing it up would provide an interesting insight. If the thrust of the example is in fact correct, then programs behind CSCs do not merely provide quantities of inexpensive energy conservation, they also provide savings spread over a whole range of other engineering improvements (see Joskow and Marron, 1992, p. 49).

Most CSCs do not start out with negative *CCE* measures, but almost all that do not start negative, start with measures that have *CCE*s very near zero. Like the negative-*CCE* regions, the zero-*CCE* beginning of most CSCs does not fit well with theory. This phenomenon might also be an artifact of omitted negative conservation measures, perhaps combined with some other selection bias, or algorithmic anomaly. For now, all that can be said is that this discrepancy between theory and practice deserves further investigation.

Now consider the example of an omitted measure that describes a partial-program. For instance, consider a program to encourage the adoption of compact fluorescents. If it attempts only modest penetration it may achieve energy savings at a cost of only 2¢/kWh, while a broader, deeper program may save twice the energy at an average cost of 6¢/kWh, still less than the cost of energy at 8¢/kWh. It may be quite tempting to those who build CSCs to include only the broader program in order to maximize the cost-effective energy savings. But if both measures

---

17. Whether this particular measure is realistic does not matter; the point is that such negative-conservation measures cannot be ruled out until they have been seriously investigated. Currently they are simply being ignored.

were included in the CSC, as they should be, it would be found that the *marginal* cost of the broad measure is 10¢/kWh, and thus only the narrow measure should be accepted.[18]

BEYOND THE STANDARD CSC

This section discusses three problems that require redefinition of the standard CSC. The problem of useful conservation measures that change the service level is solved by defining a constant-utility CSC. This definition also helps solve the problem of measures that cause rebound. Probably the most serious practical problem is the difficulty collecting accurate data. This problem is alleviated by the definition of a "technical potential" curve that is an optimistic bound to the CSC. But this solution is found to be both partial and not without danger.

The Problem of Utility and Rebound

So far we have defined a TCC curve as an isoquant of a production function of energy services. Two major problems arise when we restrict ourselves to iso-energy-service curves. First we rule out measures that provide less service than the base case but provide so much savings that the consumer is clearly better off, and second, we under-value all measures that produce a rebound effect. The first of these problems rules out a substantial portion of all proposed conservation measures, and the second undervalues nearly all services that are energy price elastic, making some appear uneconomic when they are actually very economic.

As examples of the first problem, consider that compact fluorescents turn on slowly, that efficient compressors for refrigerators make more noise, that low-flow shower heads transfer less heat at the same temperature, not to mention the differences between gas and electric stoves and problems with thicker refrigerator-wall insulation. All of these highly effective conservation measures are simply ruled out when we construct a CSC curve from an iso-energy-service curve. To get around this problem, conservationists often simply assert that the differences in service do not exist or do not matter. Since discovering the opposite would eliminate a cherished conservation measure from consideration, little if any effort is made to investigate such phenomena. The result is a serious loss of credibility and, very likely, some serious over-estimates of conservation potential.[19]

The second problem, rebound, cuts the other way. All conservation measures make energy-using equipment cheaper to operate. In many, this increases the use of the equipment. Even in cases where rebound seems impossible (who opens their refrigerator more just because it's more efficient?), there are subtle ways in which it can take place. For instance approximately 25% of all homes now have two refrigerators, the second generally being kept in the garage or basement. But widespread knowledge that old refrigerators are quite costly to run limits demand

18. For the average cost to be 6¢/kWh when the first half costs 2¢/kWh, the second half must cost 10¢/kWh.

19. Some conservation measures, such as low-E glazing, increase consumer utility. When these are included in the iso-service framework, they fail to get credit for their utility enhancements.

for them. What will happen in ten years when the retiring refrigerators cost less then half as much to run because of DOE appliance standards?

The Constant-Utility CSC

As a remedy for the first problem we introduce the concept of the constant-utility CSC as a replacement for the constant-service CSC. Recall that all conservation packages on the TCC provide the same level of energy service, which by definition means they produce the same level of consumer utility. Constant utility is the important concept and constant service is just an easy way of insuring it. So without any change in philosophy we now define the TCC curve as the lower convex bound on conservation measures that produce the same *utility* as the base case.[20] No other change will be needed in the definition of a constant-utility CSC, but we will need a method of identifying and plotting measures that are constant-utility but not constant-service.

Now consider a measure, *A*, that decreases the energy service to a consumer, and thus decreases his utility. There will be some monetary payment which would make the consumer feel fully compensated for this loss of utility. This payment may be difficult to determine in practice, but the concept itself is not vague. In fact any time it is argued that a conservation measure produces a service that is just as good but different, say the service of gas as opposed to electric stoves, the same concept of comparing the utility of different services must be invoked. This definition allows us to associate with any conservation measure, *A*, regardless of how it changes energy service, a "disutility payment", $du(A)$.[21]

Now consider the measure A^* that consists of measure *A* plus the disutility payment. By definition this new measure has the same utility as the base case. If measure *A* is described by (TCC, ΔE), then A^* is described by ($TCC+du(A)$, ΔE). Measures may now be plotted in the normal way using these “compensated” TCCs, and all the rest of our theory carries through as before.

Adjusting for the Rebound Effect

We now apply the constant-utility CSC to the evaluation of a measure that causes a rebound effect. A measure with rebound can still be described by what is sometimes called its mechanical effect; this is the amount of energy that it would save if the consumer did not react to the effective energy-price reduction caused by the increased efficiency. One reason energy-efficiency advocates may have backed away from the rebound effect is worry over the following calculation.

---

20. For residential and commercial measures consumer or employee utility is relevant. For industrial measures, profit can be substituted for utility. Thus if a conservation measure reduces profit by $x, then $x a constant-profit version of the measure would need to include a $x compensation payment.

21. du(A) must be defined relative to some base case utility level. In the case of a utility increase it will, of course, be negative.

A measure X, costing TCC, is designed so that it mechanically causes energy savings ΔME and leaves the consumer with the same utility as the base case. However when implemented, it is measured to save only $\Delta E = \Delta ME - dE$, where dE is the rebound effect, and thus its CCE is computed as $TCC / \Delta E$. But this is wrong because the realized measure is not a constant utility measure, but instead provides more utility than the hypothesized mechanical version of the measure.[22]

Recently this result has been recognized and accounted for in at least one study by deliberately ignoring the rebound effect when computing energy savings, but not when computing the *CCE* (Brown, 1993, p. 30). This method is based on the correct observation that the consumer's "takeback" cannot reduce the value of the measure, because the takeback is voluntary. In fact the first part of the takeback will provide the consumer more value than the cost of the energy, thus making her better off. As the takeback proceeds, the marginal value declines until it equals the price of energy and the process ends. By approximating the extra utility gained in this process, we can make a further correction of *CCE* for a measure with rebound.

The consumer surplus (value of utility) gained from the takeback is the area under the energy-service demand function between the original service level (S_0) and the final service level (S_1). The principal part of this is covered by $dE{\cdot}P$ (where P is the energy price) which is accounted for by simply ignoring dE in the calculation of CCE. But there is also a triangular part, just described, with an approximate area of $(P_0^S - P_1^S) \cdot (S_1 - S_0)/2$, where P^S is the price of service. By defining service, S, to equal efficiency times energy input, $f{\cdot}E$, one finds that $P^S = P/f$. These relationships are sufficient to compute that the extra consumer surplus is $\Delta CS = dE{\cdot}P{\cdot}(E_0/E_1 - 1)/2$, where E_0 is energy use before implementation, and $E_1 = E_0 - \Delta ME$. A corrected *CCE* can now be computed as $(TCC - \Delta CS)/\Delta ME$.

A more fundamental problem with rebound is that it depends on the price of energy. This makes both the CCE and the ΔE of each measure, and therefore the CSC, dependent on the price of energy. This is inconvenient because only one energy-price line can be used with a given CSC.

The Usefulness of Constructing Bounds on CSCs

From the beginning those who construct CSCs have been plagued by the problems of collecting accurate data on the *TCC* and *ΔE* of potential future conservation programs. This problem has two components: 1) determining how costly and effective the programs will be, and 2) determining the level of energy efficiency some twenty years in the future assuming the programs are not implemented. Because these problems are often viewed as simply too daunting, sometimes the choice is made to construct a "technical potential" supply curve from a "frozen efficiency baseline." "Technical potential" refers to assuming that conservation programs will achieve a 100% saturation without programmatic delays or costs. "Frozen baseline" refers to the assumption that in the no-program base case existing buildings are not retrofitted, and, for

22. See Chamberlain (1993) and Brathwait (1993).

example, that replacement equipment through the year 2010 will all be of 1990 vintage (Koomey et al., 1991).

Those choosing this strategy are well aware that “technical potential is an upper limit to the amount of efficiency that can be captured” (Koomey et al., 1991), and that a frozen baseline is a lower limit on efficiency in the absence of such programs. After a close look at CSCs reported by Lovins and EPRI, Joskow and Marron (1992) term them “explicit upper bounds on what can actually be achieved.” EPRI (undated) says the “estimated gains are technically feasible” but “practically unachievable.” Krause et al. (1988) state “technical potential represents an upper limit to savings that can be achieved”. Thus partly for reasons of simplicity a strategy is adopted of computing an optimistic bound on the true CSC. For short we will call a curve based on the “technical potential” and “frozen baseline” assumptions a TP curve. Such a curve will lie below and to the right of the true CSC, but because is puts an optimistic bound on the available cost-effective conservation, we will sometimes refer to it as an upper bound.

The looseness of these upper bounds is significant. Krause et al. (1988) put the bias of their own TP curve at nearly a factor of two. That is, they believe only about half as much conservation to be available at the price of energy as is indicated by their TP curve. Brown, a favorable reviewer, put the bias of LBL’s TP curve at a factor of 2.2, and he did not take into account the frozen baseline assumption.[23]

Optimistic bounds by their nature only allow one to conclude that cost-effective conservation *cannot exceed* some particular amount. If the TP curve indicates that conservation is not worthwhile, then one need not go to the trouble and expense of a more elaborate estimate. This is similar to the early stages of drug testing; if a drug is found to be toxic, then expensive clinical trials for efficacy can be avoided. But no amount of toxicity testing can prove the drug is efficacious, and TP curves can never prove conservation measures are cost effective.

Although their usefulness as an inexpensive preliminary screening device is substantial, this is not all that is desired of a CSC. Probably for that reason, one sometimes finds a perplexing addition to this strategy of constructing an optimistic bound.

In addition to the assumptions which make the TP curve an optimistic bound, a generally conservative approach is adopted in other aspects of the estimation procedure.[24] While this is clearly done for the sake of caution, it can undermine the purpose of an optimistic bound and may inadvertently lead to serious misinterpretation. It is desirable to put as low an upper bound as possible on the amount of available conservation, because doing so will provide the most

---

23. Brown (1993) worked closely with Koomey’s group at LBL when constructing his estimates of the “achievable potential” that corresponds to the “technical potential” reported in Koomey et al. (1991).

24. The most problematic of these assumptions from the point of threatening the consistency of the approach are those that apply to individual measures. I will not document these, because my point is only that caution should be observed to avoid compromising the definition of an upper bound. The conservative assumptions that are more likely to mislead a general reader, usually concern omitted measures and technologies. An example of these will be given shortly.

effective screen for uneconomic programs. But, if instead of simply tightening the upper bound, one makes assumptions that could generate an estimate below the true value, then one no longer has an upper bound. In this case it loses its value as a screening tool. A curve has deliberately been constructed to be an over estimate in some respects and an underestimate in others is neither an optimistic nor a pessimistic bound, and simply has no use. Whether the conservative assumptions typically introduced invalidate the TP as an upper bound, or simply make it a tighter bound is never discussed in the literature.

A second danger of introducing conservative assumptions is that they may be emphasized at the expense of the optimistic assumptions, and that non-specialist users of these studies will be given the impression that the TP curve is actually a kind of estimate that can justify the pursuit of the conservation programs described. In one case in which a TP curve was computed, the report's conclusion does not mention that the "technical potential" and "frozen baseline" assumptions are both used to produce a particularly optimistic bound on the true CSC (Koomey et al., 1991). But, it does spend one third of its time emphasizing its conservative bias: "Potentially large efficiency resources have not been included in the analysis, ..." That report's conclusion begins "This analysis has demonstrated that there are significant, cost-effective energy efficiency resources available in the U.S. residential sector" and continues "These savings represent about 40% of the frozen efficiency baseline." This gives the definite but unfortunate impression that an estimate has been made of the available cost-effective energy efficiency resources.

Since the 1991 report cited above, Koomey's group has taken account of the above considerations and expanded its methodology to include both a tighter upper bound and an true estimate of available cost-effective resources. The bound is now simply a useful step in the process of constructing the estimate, and presenting both should eliminate any chance of misunderstanding. Hopefully others who construct CSCs will follow this example. But those who still simply compute an optimistic bound should explicitly avoid any conservative or pessimistic assumptions of the type that could compromise the optimistic nature of the bound. This does not preclude efforts to make the bound as tight as possible.

CONCLUSION

Unlike a true supply curve, a conservation "supply" curve (CSC) describes the cost of conserving an input of production with output held constant. Using this definition, a CSC can be derived from a production function and the prices of the other inputs. This process reveals that CSCs do not stop at the origin, but continue to the left into the region of negative conservation, and that they are monotonically increasing. It also shows they will only take on negative for levels of energy consumption so high that the cost of removing waste energy becomes a dominant factor.

The interaction of conservation measures has been recognized and accounted for since the begining, but the algorithm used to account for these interactions produces suboptimal CSCs. An optimal algorithm reveals the need to include negative conservation measures and partial measures when possible. Omitting either type of measure can lead to more conservation than is optimal.

The basic CSC definition assumes a constant production level for energy service, thereby ruling out all conservation measures that change the service level. This eliminates many of the best conservation measures. This problem can be overcome by basing a CSC on constant utility. One consequence of this approach is that measures which cause rebound are seen to fair much better than they do under standard analysis.

One of the most vexing problems of CSC construction is the difficulty in obtaining unbiased estimates of the required data. This has led to the frequent use of "technical potential" curves which are optimistic bounds on CSCs. These are found to be useful for screening, but potentially misleading when not explicitly distinguished from estimates. The appendix solves the problem of aggregating cost and energy savings over time, and shows that standard CSCs cannot be constructed in a world with changing energy prices. Two alternatives are proposed.

These observations caste serious doubt on conservation curves that begin with negative marginal costs, and rule out those that are not monotonic. This should prompt the correction of some well known CSCs. The tendency for all non-negative CSCs to intersect the origin, is seen to be puzzling and to derserve further investigation. In spite of these questions, the main contribution of this paper should be to help legitimize CSCs by positioning them within the standard economic framework, and to build a much firmer theoretical base for their construction and use. If careful attention is paid both to the theory of construction and to the practical questions of utility, rebound and upper bounds, then CSCs can be used for their intended purpose.

APPENDIX

Summary

This appendix will show that the traditional formula for CCE is incorrect, though not by enough to matter when discounts rates are low. The significance of this fact lies in directing us to a correct formulas which is actually far more general and just as simple as the traditional formula. The correct formula is:

$$CCE = TCC / \Delta E, \quad \text{where}$$

$$TCC = \mathrm{PV}(c(t)), \quad \text{and}$$

$$\Delta E = \mathrm{PV}(e(t))$$

where *PV*(.) is the present value operator defined by $\int_0^\infty (\,.\,)\, e^{-r \cdot t} dt$, $c(t)$ is the stream of measure costs including initial cost, operation and maintenance, $e(t)$ is the flow of energy savings, and r is the continuous discount rate. These formulas cover all cases so long as the price of energy is constant. In the case of fluctuating energy prices, this definition cannot be used to build a functioning CSC. Several of the ideas reported in the appendix were independently developed by Norris (1994).[25] The last section of the appendix, examins the questions for negative CSC, free disposal and demand charges for energy capacity.

The Traditional Definition of CCE

The original method of constructing a CSC assumes that the measure cost, C, occurs entirely at time zero, and that energy is saved at the constant rate of ΔE/year. In order to make the cost and savings more comparable, cost is "annualized" according to the formula

$$\text{Annualized Cost} = AC = \frac{C \cdot d}{1 - (1+d)^{-n}}$$

where d is the consumer's annual discount rate and n should be the number of years of energy savings. CCE is then defined as AC/ΔE.

The variable n was defined as the time in which the investor "wants to recover the investment" (Meier, et al., 1983), then as "the number of years over which the investment is written off or amortized" (Rosenfeld, 1993), though Koomey et al. (1991) define it correctly as "the lifetime of the conservation measure". The confusion probably arises from the origin of the annualizing formula, which is based on a "capital recovery factor." Its formula is designed to give the annual, year-end payment that would be necessary to repay and initial investment over

25. Norris writes "A mathematically equivalent definition of CCE which allows for treatment of dynamic and uncertain energy savings is 'the cost of the investment divided by the present value of the measure's lifetime energy savings'." He also tackles the problem of non-constant energy price and finds it can be solved by using the following capital recovery factor when computing CCE: crf = $(1/E_0 \cdot P_0) \int P(t)E(t)e^{-dt}\, dt$. This is probably equivalent to the solutions given at the end of this appendix.

a "write off" period of *n* years. Some have supposed that the annual periodicity of this payment is useful because it mimics the periodicity of energy bills. But this style of thinking is completely off the mark.

The problem is to find a payment stream that is exactly proportional to the energy savings stream and with the same value as the actual payment stream. With two proportional streams we will always get the same answer no matter when we divide cost by energy savings, so there can be no doubt about the unit cost of this energy. For example if the $c(t)$ = 5¢/hour × $e(t)$, where $e(t)$ is measured in kWs, then *CCE* must be 5¢/kWh. Since the traditional definition of CCE assumes constant energy savings, the appropriate payment stream would be a continuous constant payment stream. This is not badly approximated, for small discount rates, by a payment stream that occurs in annual lumps, so the traditional formula is quite close provided *n* is taken to be the lifetime of the energy savings stream. Using a continuous level payment stream, the correct formula is:

$$\text{Levelized Cost} = LC = \frac{C \cdot r}{1 - e^{-r \cdot n}}$$

where *r* is the consumer's continuous annual discount rate. CCE is then defined as LC/ΔE.

This formula is a first step towards a far more general formula. The second step generalizes to measures whose cost does not all accrue at time zero. Since paying the PV(c(t)) at time zero is equivalent to paying c(t), the levelization formula for a fluctuating cost is

$$LC = \frac{PV(c(t)) \cdot r}{1 - e^{-r \cdot n}}$$

and the correct formula for CCE is still LC / ΔE. The next necessary generalization allows the flow of energy conservation to be non-constant.

Using the Present Value of a Physical Energy Flow

The present value method of computing CCE is simple and is completely general in a constant energy-price world. We derive it in two ways. Again we assume a cost stream, $c(t)$, and a conserved energy stream, $e(t)$.

The first approach is to find a cost stream with the same present value as c(t) that is exactly proportional to the energy conservation stream. Having done this we compute CCE(t) = c(t)/e(t), which is the instantaneous price of conserved energy at all times. Since c(t) has been constructed to be proportional to e(t), CCE(t) will be constant. Since it is constant it must equal the average CCE, no mater how we weight different times. A little algebra shows that such a cost stream, c(t) is given by [PV(c) / PV(e)]·e(t). Thus the general formula for CCE is:

$$CCE = \mathrm{PV}(c)/\mathrm{PV}(e)\,.$$

The second approach is through the definition of the cost of conserved energy. Define CCE to be the value which if greater than the avoided price of electricity indicates that a conservation project is uneconomic and if lower indicates it is economic (Meier, 1982, p.45). Now a measure is economic if and only if PV(c) < PV(ac), where c(t) is the cost of the conservation measure and $ac(t)$ is the avoided cost of the conserved energy. If the avoided cost stream is simply the conserved energy stream times a constant avoided price, i.e. if $ac(t) = P \cdot e(t)$, then we can factor the avoided price out of the present value operator and find:

$$\mathrm{PV}(c) < P \cdot \mathrm{PV}(e)$$

$$\text{or} \qquad \mathrm{PV}(c)/\mathrm{PV}(e) < P$$

Thus if the avoided price, P, of electricity is constant, then conservation projects can be evaluated simply by comparing CCE and P if and only if CCE is defined as PV(c)/PV(e). Thus the definition of CCE is again:

$$CCE = \mathrm{PV}(c)/\mathrm{PV}(e)$$

Solving the General Problem: Fluctuating Energy Price

The second derivation of *CCE* required an assumption of constant avoided energy price, *P*, which was not required by the first definition. That is because the second approach, but not the first, requires *CCE* to be capable of separating economic from uneconomic measure using the price of energy as a comparison. This points to a serious problem with even this definition of *CCE*. As defined, CCE cannot be used to compare the economics of two conservation projects unless avoided energy price is constant. As the first approach shows, this does not mean that CCE is not well defined in such a world; it just does not have its most essential property. (The problem arises because a low-CCE measure may save energy when its price is high, while the high-CCE measure may save energy when the price is low.) Thus projects cannot be ranked by their CCEs in a world with non-constant avoided energy price. This may be problematic since Meier (1982, p.45) tells us that “Independence [of CCE] from energy prices is particularly desirable when the most volatile element of the conservation investment decision is energy price.” A definition of CCE that does not allow project comparisons does not allow the construction of a useful CSC. There are two solutions to this problem.

The first is to redefine *CCE* as follows:

$$CCE = \frac{P_0 \cdot \mathrm{PV}(C)}{\mathrm{PV}(P(t) \cdot dE(t))},$$

and then proceed as usual. Unfortunately this does make CCE dependent on the price of energy. The second solution, suggested by Carl Blumstein, uses PV($P(t) \cdot dE(t)$) in place of ΔE when constructing the TCC curve. All points on this curve and below a 45° diagonal through the origin would be economical. Of course, since the present value of a conservation measure is just PV($P(t) \cdot dE(t)$) − TCC, this is just the standard rule of accepting any investment that has positive present value.

Negative CSCs When Energy Disposal is Not Free

This discussion concerns the question of whether the base case is technically efficient, and thus the question of how much savings should be attributed to energy conservation and how much to non-energy efficiency gains. It does not call into question the total savings of a measure relative to the base case. The problem of differing peak and off-peak prices is adressed incidently.

A well know and elementary result of the theory of the firm, states that if the “hypothesis of ‘free disposal’” (Varian, 1978, p. 6) is true, then a production isoquant is monotonic. In our case this means that if energy can be disposed of in the production process without cost, then the TCC will be monotonic, and the CSC will be non-negative everywhere.

Generally, electricity can be disposed of very cheaply by means of a resistance heater, so one might think the hypothesis is nearly satisfied.[26] But, true though this may be, the question of free disposal is much more complex.[27] First the waste heat may need to be disposed of at some expense, either because it is annoying to the consumer or because it affects the operation of the appliance, for instance a refrigerator.

But the second objection to free disposal is the most subtle and interesting. Particularly when constructing a regional supply curve, one may wish to investigate not the conservation of electrical energy, but of the chemical energy used to produce electricity. This means that the generating plants, and in fact the entire electric utility must be considered as part of the production process. This being the case, the cost of the electricity is seen to be partly a payment for chemical energy (fuel), and partly a payment for non-energy inputs: generators, transmission lines and even labor. Often this distinction is at least partially made in the billing through the use of demand (capacity) charges, and sometimes reductions in these demand charges have correctly been subtracted from other non-energy costs in the calculation of CCEs (Piette et al., 1989). Since all (chemical) energy use is accompanied by at least some non-energy costs (even off-peak there are energy related maintenance costs), there is no free disposal of chemical energy in an electrically powered production process, such as an electric appliance. In fact for commercial lighting, which operates on peak, this effect is quite strong. (Although energy could still be cheaply disposed off peak, this would require the expense of a timing device, and as noted would still not completely solve the problem.)

Thus we conclude that even considered from the most rigorous economic perspective, the hypothesis of free disposal if often quite false for a chemical-energy input, and thus negative CSCs are possible. Of course such CSCs must be used in conjuntion with the price of chemical energy. This is not meant to imply that actual negative CSCs were correctly constructed. Until

---

26. Tthis insight was provided by Carl Blumstein.

27. A difficult position to defend, but one that should not be ignored without careful consideration, is that while it is technically possible to produce an electric appliance with a wasteful heater, this option would not be permitted by the sociology of the market. Accepting this position, implies that some technically inefficient appliances may be effectively on an upward sloping tail of the production isoquant.

the correct construction algorithm is used and a full investigation is made of possible low- and negative-conservation measures, this will remain an open question.

REFERENCES

Blumstein, Carl, and S. E. Stoft (1995). *Technical Efficiency, Production Functions and Conservation Supply Curves*, POWER Working Paper, PWP-029, University of California Energy Institute, April.

Brathwait, Steven D., D. Caves and P. Hanser (1993). “The Complete and Unabridged Measure of DSM Net Benefits: What We’ve Been Missing.” Presented at 6th National Demand-Side Management Conference, published by the Electric Power Research Institute.

Brown, Richard (1993). Estimates of the Achievable Potential for Electricity Efficiency Improvements in U.S. Residences, masters project for the Energy and Resources Group, University of California, Berkeley, May 18.

Chamberlain, John N. and P. Herman (1993). “Why All “Good” Economists Reject the RIM Test.” Presented at the 6th National Demand-Side Management Conference, published by the Electric Power Research Institute.

Electric Power Research Institute (undated). *End-Use Energy Efficiency*. Palo Alto, CA.

Electric Power Research Institute (1990). *Efficient Electricity Use: Estimates of Maximum Energy Savings*, EPRI CU-6746, prepared by Barakat & Chamberlin, Inc. Palo Alto, CA. March.

Goldstein D., R. Mowris, B. Davis, and K. Dolan (1990). *Initiating Least-Cost Planning in California: Preliminary Methodology and Analysis*, Natural Resources Defense Council and the Sierra Club, prepared for the California Energy Commission, Docket No. 88-ER-8, Revised May 10.

Huntington, Hillard G. (1994). “Been top down so long it looks like bottom up to me,” *Energy Policy*, Vol 22, No 10: 833-838.

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

Koomey, J. G., C. Atkinson, A. Meier, J. McMahon, S. Boghosian, B. Atkinson, I. Turiel, M. Levine, Bruce Nordman, and P. Chan (1991). *The Potential for Electricity Efficiency Improvements in the U.S. Residential Sector*. Lawrence Berkeley Laboratory. LBL-30477, July.

Lovins, Amory (1985). "Saving Gigabucks with Negawatts," *Public Utilities Fortnightly*, March 21.

Lovins, A. B., R. Sardinsky, P. Kiernan, T. Flanigan, B. Bancroft, and M. M. Shepard (1986) *Competitek*, Rocky Mountain Institute, Snowmass, Colo..

Lovins, Amory B. and L. Hunter Lovins (1991). “Least-Cost Climatic Stabilization”, *Annual Review of Energy and the Environment* 16:433-531.

Meier, Alan K. (1982). *Supply Curves of Conserved Energy*. Lawrence Berkeley Laboratory. LBL-14686, May.

Meier, A., J. Wright and A. H. Rosenfeld (1983). *Supplying Energy Through Greater Efficiency*, University of California Press, Berkeley, CA.

New England Energy Policy Council (1987) Power to Spare: *A Plan for Increasing New England’s Competitiveness Through Energy Efficiency*, July.

Norris, Gregory A. (1994). “Energy Conservation Potential Uncertainty Analysis”, mimeo, 16 February.

Northwest Power Planning Council (1991). 1991 *Northwest Conservation and Electric Power Plan Volume I*, 91-04, April.

Panel on Policy Implications of Greenhouse Warming (1992). *Policy Implications of Greenhouse Warming*, pp 171-200, National Academy Press, Washington, DC.

Piette, M. A., F. Krause, and R. Verderber (1989). *Technology Assessment: Energy-Efficient Commercial Lighting*. Lawrence Berkeley Laboratory. LBL-27032, March.

Rosenfeld, A. H., et al. (1993). "Conserved Energy Supply Curves for US Buildings," *Contemporary Policy Issues*, (11), 45-68.

Rubin, E. S., et al. (1992). "Realistic Mitigation Options for Global Warming," *Science*, (257),148-149,261-266.

Ruff, Larry E. (1988). "Least-cost Planning and Demand-side Management: Six Common Fallacies and One Simple Truth," *Public Utilities Fortnightly*, 19, April 28.

Usibelli, A., B. Gardiner, W. Luhrsen, and A. Meier (1983). *A Residential Conservation Data Base for the Pacific Northwest*. Submitted to: Bonneville Power Administration, Lawrence Berkeley Laboratory. LBL-17055, November.

Varian, Hal R. (1978). *Microeconomic Analysis*, New York, W. W. Norton & Company.