BEEN TOP DOWN SO LONG IT LOOKS LIKE BOTTOM UP

Hillard G. Huntington

EMF OP 41

May 1994

Energy Modeling Forum
Stanford University
Stanford, CA 94305-4022

Phone: (415) 723-0645
Fax: (415) 723-4107
Email: HILLH@FORSYTHE.STANFORD.EDU

ABSTRACT

While technologists emphasize suboptimal choices by individuals and organizations, many economists have ignored behavior exhibiting technical inefficiency. This paper selectively discusses efforts by some economists to identify and measure technical inefficiency. As applied to the energy efficiency "gap," "bottom up" economics would seek to measure significant departures from best-practice technology, distinguish between inefficiency that is more (less) wasteful of energy than other inputs, and focus on developing testable hypotheses about technical inefficiency. Support for widespread intervention in energy use decisions may not necessarily follow from such an inquiry.

BEEN TOP DOWN SO LONG IT LOOKS LIKE BOTTOM UP[1]

Hillard G. Huntington[2]

If an economist were to stumble across a technical inefficiency[3] by accident, could he identify and measure it as such? If decisionmaker "inertia" and other related inefficiencies are excluded by assumption rather than by empirical measurement, can there be a meaningful discussion between economists and technologists on the apparent underinvestment in energy efficiency or the "gap"? This challenge, raised by a participant in a recent Energy Modeling Forum (EMF) study on energy conservation, serves as a focal point for this paper.

Simply raising this issue can create misperceptions about the nature of the energy efficiency debate. First, the empirical burden of proof in the "gap" debate by no means rests solely with economists. One could just as easily pose the counter question: why don't technologists empirically test whether and by how much the specific market failures actually do influence energy efficiency decisions?

Second, addressing the technical inefficiency issue should not detract from efforts to use standard economic analysis to help define the issues. Jaffe and Stavins (1994) provide a much-needed framework for clarifying key terms (e.g., barriers, failures, and alternative concepts of the conservation potential), Hassett and Metcalf (1993) develop an economic rationale for what appear to be unreasonably high discount rates, and Nichols (1994) shows how traditional economic analysis can be used to indicate the degree of inefficiency required (which is large by his estimates) to justify a particular demand-side management program.

Third, advocates of more-aggressive intervention in energy use decisions may find that an inquiry into the extent and causes of technical inefficiency does not lend the unambiguous support for the policy prescriptions that they frequently offer. It may be that the economist's policy advice of minor modifications in market institutions is preferable to widespread end-use standards or utility demand-side management programs in many circumstances.

Despite these caveats, I think that economists might usefully explore the inefficiency issue as it relates to the energy-efficiency "gap". Not only would this research provide better estimates on whether such inefficiencies are large or small, but it could also build the foundations for discussing technical inefficiency issues related to the "gap" in more precise terms.

In this paper, I selectively review some empirical studies by economists that appear to find evidence of technical inefficiency or of behavior that would be considered anomalous from the perspective of traditional economic theory. I draw first upon some studies from experimental economics, which test participants' responses to different incentives and rules under laboratory conditions, and then proceed to studies based upon real industry or sector data.

Such an inquiry might be labeled "bottom up" economics. In the debate about the energy efficiency gap, technologists have argued that significant unexploited opportunities exist for cost-effective investments in energy efficiency. In support of this position, they have adopted a "bottom up" perspective that relies heavily on cost and performance characteristics of specific technologies and on the concept that many decisionmakers might be behaving suboptimally. This approach contrasts with a "top down" approach, favored by many economists more skeptical of these opportunities, that focuses upon general market forces influencing efficient decisionmakers in their use of energy.

TOWARDS A TESTABLE HYPOTHESIS

Economists embrace the efficiency assumption because it allows them to develop testable hypotheses from the least complex of frameworks. For example, from an analysis of the consumer maximizing his utility subject to an income (and time) constraint, one can develop a testable hypothesis about how a change in energy prices will affect energy consumption. The test is cast in terms of a null hypothesis--one that the researcher wishes to reject thereby allowing him to accept the alternative hypothesis. In this example, the null hypothesis is that energy prices have no effect on energy consumption. Across a range of different consumers, regions, and countries using different specifications of the demand relationship, researchers have been able to reject this hypothesis based upon empirically observed data. Hence, the widespread belief among economists that energy prices have an effect (often an important one) on energy consumption. The efficiency assumption is never tested directly but rather serves as a vehicle for testing other hypotheses.

Suppose now that attention is turned from developing testable hypotheses about the effect of prices and other factors on individual or aggregate behavior towards promoting a dialogue with those more skeptical of the efficiency concept (hereafter, called "technologists", an imprecise term at best). One could envision a null hypothesis that consumers are inefficient (in the sense discussed below) in their decisions about energy as well as other products. As structured, this hypothesis has at least a chance of not being rejected when tested with observable data, but it does require some precise measures of inefficiency. The key point is that the hypothesis should precede the observation (ex ante analysis) rather than the reverse (ex post analysis).

Such inefficiency need not result from poor decision-making skills, as is sometimes implied by the term, technical inefficiency. Liebenstein (1966) discusses a wide range of contractual or organizational limits (which he calls X-inefficiency) that prevent firms from acting like a single, efficient decision making unit. He argues that the losses due to the more systemic problem of X-inefficiency dominate the relatively minimal welfare losses resulting from the "allocative" inefficiencies studied by most economists. Recently, DeCanio (1993) has shown how conflicting incentives within a firm create organizational "slack" when firms make decisions about energy efficiency. Many economists have developed ex-post explanations that make observed X-inefficiency anomalies conform to hypothesized rational behavior, but few have offered ex-ante predictions that could be either confirmed or refuted through empirical tests (Frantz, 1990).

IMPUTED DISCOUNT RATES AS EVIDENCE OF INEFFICIENCY

The most commonly cited evidence for consumer inefficiency is unfortunately not a clear barometer for measuring its presence or absence. Energy consumers appear reluctant to pay higher prices for equipment that will lead to relatively large energy savings and reduced fuel costs within a short span of years. By comparing the future energy cost savings with the initial equipment costs, one can compute an implied required rate of return that would represent a breakeven point for the investment. These required rates of return range from 20% (Hausman 1979) to much higher (Train 1985), making it appear that consumers are requiring excessively high rates of return to invest in energy efficiency.

Economists have noted some important shortcomings in many attempts to infer inefficient behavior from such estimates: measurement errors in equipment costs, energy savings, or both; the existence of certain omitted transaction costs that private consumers should bear; the presence of other amenities of value by the consumer; or the need for compensation for risk (Sutherland 1991, Jaffe and Stavins 1994, and Nichols 1994).[4] In addition, such imputed rates ignore the opportunity cost of acting today rather than delaying the decision in hopes of resolving some important uncertainties with equipment that cannot be easily scrapped (Hassett and Metcalf, 1993).[5] Home insulation represents a good example; the consumer does not have to worry about others competing for the opportunity to invest in his home, delay may (but not necessarily) improve his knowledge about the investment, and used insulation can not be removed easily and sold if the investment turns sour.

ECONOMIC BEHAVIOR IN THE LABORATORY

Another source of information about consumer choices consists of experiments by economists to analyze how people respond to incentives in different market environments. The researchers conduct these experiments in an economics laboratory environment where they can carefully control the costs, payoffs, and other rules of the game that might mimic different market settings. Although there are concerns about whether the participants (usually university students) really have the proper incentive or are representative of real-world decisionmakers, the researchers find results that are remarkably robust across different populations and conditions.

These studies have confirmed some important economic principles as well as raised some intriguing puzzles. Although participants do not purposely maximize their monetary earnings or well-being, their repeated interactions in buy-and-sell situations often lead them to economically efficient outcomes (Smith, 1994). Moreover, they often do so with limited information, far less than the perfect information assumed by some economic theorists. However, institutions matter a great deal. Altering the rules under which trades are transacted (e.g., different types of auctions) changes available information and incentives for individuals, thereby leading to different outcomes.

Thaler (1993) summarizes some puzzling results that emerge from such studies. First, context is extremely important. As one example, participants require a higher price to relinquish something they own than the cost of obtaining a new one. This disparity conflicts with traditional economic analysis in which a consumer's willingness to exchange between goods is invariant to which good he owns. Theorists have resorted to psychological explanations--an endowments effect produced by an aversion to a loss--to explain such behavior.

Second, participants express a strong preference for the status quo or their reference position. Once familiar with a product they like, consumers appear reluctant to change.

Experience appears to matter, thus explaining the widespread use of rebates and coupons in certain markets to attract initial customers. While not technically an inefficiency, the experience factor emphasizes the need for a more dynamic approach to consumer choice rather than the more static analyses frequently employed in economics.

Third, participants appear to require higher rates of return for shorter payback periods than for longer ones and for smaller investment amounts than for larger ones. Participants are influenced by the magnitude of the investment and payouts as well as by how long they are required to wait to realize the gains (or losses). These experiments eliminated both transaction costs and risks, two factors that could help to explain the anomalies in real-market settings. Transaction costs refer to additional expenses or lost opportunities that decisionmakers incur to buy or sell a product but that are often unmeasured. They are likely to represent a larger share of total costs for smaller investments and their negative effect on the investment return more dramatic for shorter investment periods. Similarly, uncertainty about equipment costs could create greater risks for investments with shorter horizons.[6] Being able to isolate and remove the effects of both these influences represents an important advantage of the experimental economics approach.

These results show that NPV estimates of the returns to energy-efficient equipment, which indicate the highest required rates of return for small appliances with payback periods of a few years, are part of a more general behavioral phenomenon rather than an example of how energy use differs from other decisions. But the empirically-observed time variance of discount rates opens a number of messy issues for economists, because it suggests that there may be opportunities for profit-seekers to borrow funds at low rates over long periods and to loan them at higher rates over shorter periods. Many of the "top-down" models used to evaluate climate change policy over decades exclude such market imbalances and require a constant discount rate.

It would appear that a richer model of intertemporal choice would result from further investigation of this apparent anomaly.

It should be emphasized that the above anomalies could permeate many different types of behavior, not simply those decisions about energy use. They do not necessarily imply overuse of energy and therefore do not justify, on their own, public intervention to reduce energy consumption. For example, it appears likely that some consumer inefficiency prevails in markets where purchases occur infrequently, but energy-using equipment represents one of many such items that households and commercial establishments buy only occasionally. Whether these consumers are purchasing too much or too little energy will depend upon all the goods and services that they buy. Thus, the relevant criteria is how these anomalies influence all markets, not simply the markets for energy-efficient equipment.

Upon reviewing the evidence from many studies indicating apparent anomalous behavior, Thaler (1992) concludes that many current economic models reveal "how people ought to behave" to satisfy certain objectives, not "how people actually do behave." Rather than rejecting economic theory outright, however, he advocates that it be used as a foundation to build more robust behavioral models incorporating factors that appear important in many of the experimental studies of people's behavior. In a similar way, studies that measure technical inefficiency in actual markets or sectors may have the same effect, in addition to providing a much needed conceptual structure for addressing the largely ignored question in the "gap" debate: how big and widespread are the apparent inefficiencies?

TECHNICAL INEFFICIENCY

Measuring inefficiency from real-world data is a tricky proposition. Economists differentiate between technical (considered here) and allocative (considered briefly below) efficiency.

Technical efficiency for the economist is tantamount to best practices (BP) for the technologist. In the economic paradigm, energy consumers and firms would always use BP unless it were too costly to convert from current practices.

The curve and the dashed line in Figure 1 represent technical inefficiency. Each firm in the industry uses energy inputs (measured along the horizontal axis) and other inputs (along the vertical axis) to produce a standardized product. Along the BP frontier (labeled in the figure), any firm is using the fewest inputs to produce a given amount of the final product.[7] The mix of inputs can change along this curve without affecting output, thus allowing a firm to use more energy and less of other inputs if energy should be relatively inexpensive.

Now suppose that one observes a firm operating at point A in the diagram. It uses more of both energy and other inputs to produce the same amount of the good than would be required by the BP. Its use of both inputs could be reduced proportionally without reducing output until the firm reached a point B lying along the BP frontier. A standard approach for measuring the degree of technical efficiency is to compute the percentage of observed inputs that are required to maintain the observed output, or OB/OA.

This measure of technical inefficiency assumes that the firm at point A is equally inefficient in its use of energy and other inputs. As the firm moves back along the ray OA, it reduces its use of all inputs equally. Conserving energy alone, holding other inputs constant, would shift the firm in a leftward, horizontal direction away from point A, but it may be more cost effective to reduce usage of the other input more. Therefore, policies aimed exclusively at saving energy do not necessarily overcome the problem of technical inefficiency.

In this approach, the ability to identify the BP frontier is critical. Past studies have used both linear programming and statistical techniques to determine the location of the frontier in particular industries. In the programming approach, a set of optimization constraints is used. In

Figure 1.

Measurement of Inefficiency

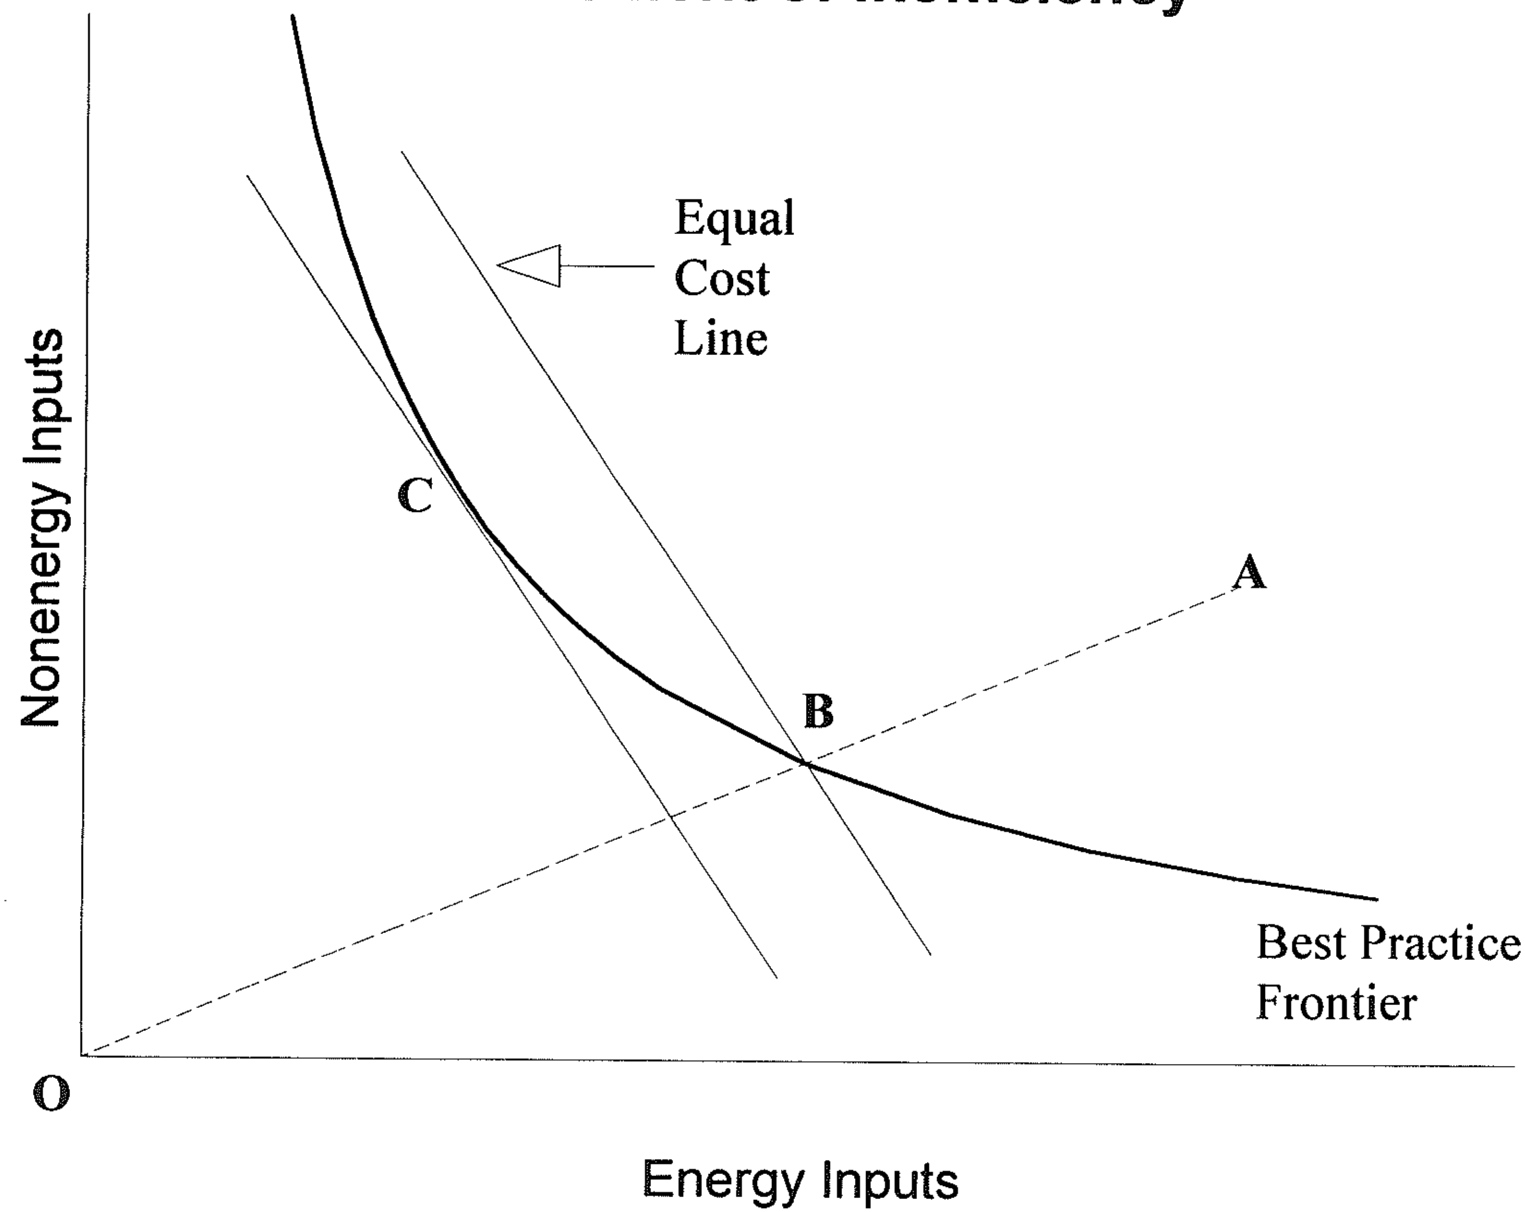

the statistical approach, the data is fitted to an equation relating physical output to a set of physical inputs. In addition to the standard randomly distributed error term, the econometric approach specifies a second error term with a skewed distribution to allow observations to lie above the frontier (inefficiency) but not below it.

Past studies have applied these techniques to various manufacturing industries, electric power generation, airlines, and agriculture in a number of different countries. Estimated efficiency varies not only with the industry or country, but also with whether the researchers used the deterministic programming or the stochastic approach and with the form of the assumed underlying production function. Caves (1992) cites four separate studies using similar methodologies that estimated efficiencies of 62.5% for Brazilian manufacturing, 55% for the Colombian apparel and footwear industry, 62-68% for Indonesian weaving, and 71-94% for French manufacturing. He concludes from these estimates that enough technical inefficiency exists to justify further investigation of this issue.

Of direct interest here are several recent studies focusing upon energy use. Ferrier and Hirschberg (1992) found a very high level of efficiency (93.9%) in controlling climate (temperature) in a sample of 65 public office buildings in Minnesota. In their deterministic linear programming formulation, outputs include the products of: (i) cooling degree days and area cooled; (2) heating degree days and area heated; (3) cooling degree days of occupied areas and number of occupants; and (4) heating degree days of occupied areas and number of occupants. Inputs include electric and non-electric energy use, total insulated surface area (walls, roof, and glazed windows), and ratio of floor to roof area.

Using a similar programming approach, Boyd et al (1993) found mean efficiencies for most years in the 1972-88 period that were above 80% among steel plants specifically classified as minimills using electric arc furnaces. Efficiency tends to be higher in prosperous years and lower during recessions. They concluded that electricity use could be reduced by from 8 to 34% (with the higher reduction associated with recessions) if all firms held output constant while operating along the BP frontier. In extensions to other industries, Boyd et al (1992) found comparable mean levels of technical efficiency in hydraulic cement (77%), blast furnaces and steel mills (83%), and paper mills (84%) and higher levels in nitrogenous fertilizers (93%), primary aluminum (94%), and alkalies and chlorine (99%). Adoption of best practice production would on average reduce energy use by 20.2%, 23.6%, 18.5%, 8.4%, 6.0% and 0.2%, respectively, in these six industries, again holding output constant.

Of course, moving to the BP frontier would allow the industry to be more productive, increasing output. It would also involve proportionally similar savings in other inputs, some of which may be even more valuable to use more efficiently than energy. Thus, while providing important information about possible inefficiency in this industry, their estimates cannot be viewed as actual energy savings or as justification for mandated energy conservation.

This conclusion becomes even stronger when they try to explain why certain steel plants are more inefficient than others. The dominant factors were how the aggregate economy affected capacity utilization (discussed above) and the vintage of the plant's capital stock. Plants with newer vintages are more technically efficient. Thus, a more effective policy might be to correct for any tax distortions biasing investment decisions away from building new plants.

Capital stock vintage often emerges as an important factor in explaining differences in technical inefficiency across different units in the studies surveyed by Caves (1992). Other factors that are sometimes important in these studies include plant size, high seller concentration, and tariff protection.[8] Of direct relevance to the energy efficiency "gap" issue, information stocks were important in studies of agriculture in both the U.S.A. and developing countries.

There are some important limitations in the empirical application of this methodology. First, attempts to compare efficiencies across industries within one country (Caves and Barton, 1990) or across countries (Caves, 1992) fail to find robust conclusions about where inefficiency may be more prevalent or what factors cause these differences. Indeed, studies of Soviet enterprises conducted during the late 1970s and the 1980s (surveyed by Murrell (1991)) appear to have found levels of inefficiencies no greater than those in Western industrialized economies--an incongruous result, at best. These problems underscore the need for further refinements to make the approach more useful for policy issues. Second, the approach assumes that inefficient firms are equally wasteful (proportionally) of all inputs, while many analysts espousing the "bottom up" view think that consumers and firms are more wasteful of energy than of other inputs. It would be important to test this contention. If refuted, it would provide policymakers with little reason to isolate on energy use to the exclusion of other key inputs.

While measures of technical inefficiency are of interest, the methodology's most important contribution may be to help define terms and concepts more precisely than in the past discussions of the "gap." The methodology rests upon a mathematical and quantitative structure for defining and estimating production and cost relationships that would provide analysts with the tools to refine ambiguous concepts and seek common ground on technical points.

ALLOCATIVE INEFFICIENCY

Although research on measuring technical inefficiency is growing rapidly, economists have traditionally devoted much more attention to allocative inefficiency. A firm could be on its BP frontier, operating with the best-practice technology, but still be economically inefficient. Allocative efficiency requires, among other things, that the firm gets the input mix exactly right for minimizing costs, given the input prices.

The framework for measuring technical inefficiency represented in Figure 1 can be expanded to incorporate the firm's allocative inefficiency if input prices are readily available. The solid diagonal lines are equal-cost lines, along which inputs can be substituted, given relative input prices, without changing total production costs. In the discussion of technical inefficiency, the firm moved from point A to B, eliminating wasteful input use, and reducing total costs by shifting its cost line from A (not shown) to B. Once on the BP frontier, however, the firm can do even better, given the prevailing input prices. It can reduce energy use, employ more nonenergy inputs, and move its total cost curve in toward the origin until it reaches point C, where it is tangent to the BP frontier. The slope of this line does not change as long as relative input prices remain fixed. Given the relative input prices, this will be the only point along the BP frontier that is allocatively as well as technically efficient for the firm. The shift from points B to C removes the firm's allocative inefficiency.

Many empirical economic studies of firm behavior assume allocative as well as technical efficiency. Researchers frequently estimate a cost function under the assumption that all firms are price-takers rather than price-setters and are minimizing their costs. This is an empirical convenience, because it is difficult to obtain robust results from direct estimation of the production function.

However, this assumption need not always be maintained. Boskin and Lau (1992) provide an interesting recent example where they attempt to remove many of the traditionally maintained assumptions about technical progress, economies of scale, and allocative efficiency in order to evaluate the magnitude of productivity growth and the direction of technical progress (i.e., which inputs are used more or less). They estimate a "meta" production function at the aggregate economy level for five major OECD countries. They assume that all countries use the same general production function but that the quality of measured inputs can vary across countries and over time. In their estimates, they allow the data to determine whether inputs are paid a price equal to their marginal product and find that labor may be slightly overpaid (presumably due to market power) in the United States and United Kingdom. Thus, it is conceptually possible to estimate the extent of allocative inefficiency, although this approach is not the norm in empirical economics.

POLICY IMPLICATIONS

That researchers have found some degree of technical inefficiency in various industries is not too surprising. But how much is too much? The 93.9% level of technical efficiency achieved within public office buildings in Minnesota appears to be quite high, while the 55% discovered for Colombian apparel and footwear industry appears quite low. Between these two estimates, there appears to be a wide range of possibilities that will alarm some analysts and policymakers but not others.

Moreover, even if most researchers agree that inefficiency is widespread, the policy implications are far from clear. One needs to know the causes of inefficiency before appropriate policies can be designed to treat the root causes of the problem rather than its symptoms.

For example, suppose that strong consumer myopia or imperfections in the capital market prevent homeowners from making many investments. Should government or the utility provide the funds but restrict it for energy-efficiency improvements only? The homeowner might be better served by having access to the funds to be spent on items of most importance to him. If he decides to repaint his home's exterior rather than to invest in a new roof, is he making an unwise decision?[9] A roof expert may tell him that a new roof is the best investment he could possibly make to keep his home dry, but he may need a new exterior coat of paint more. But the roof expert, who controls the loan, prevails, and next spring the homeowner finds the interior of his home under attack from moisture seeping through the walls.

CONCLUSION

Technical inefficiency in consumer decisions and economic activities appears sufficiently pervasive as to make "bottom up" economics an interesting and useful line of inquiry. Acceptance of several of the major points made here leads to expanding the economic research agenda to include several issues, including:

- Can significant departures from best-practice technology (i.e., technical inefficiency) be identified and measured meaningfully?
- Can energy-using inefficiency that is more wasteful of energy than of other inputs be distinguished from neutral or even energy-saving inefficiency? and
- Can ex ante tests of inefficient behavior be developed prior to observation rather than relying solely upon ex post explanations of apparent anomalies already observed?

Investigation of these issues should not prejudge policy prescriptions. In particular, a careful evaluation of these issues may not justify the "bottom-upper's" faith in widespread government intervention in the markets for energy and energy-efficient equipment. Government failures create economic distortions just as do market failures. At the same time, economists pursuing these issues might find institutional changes that influence how organizations respond to incentives or that alter how prices and information are transmitted in these markets.

REFERENCES

Bogle, John C., *Bogle on Mutual Funds: New Perspectives for the Intelligent Investor*, New York:Irwin, 1994.

Boskin, Michael J. and Lawrence J. Lau, "Capital, Technology, and Economic Growth," in Nathan Rosenberg, Ralph Landau, and David C. Mowery, editors, *Technology and the Wealth of Nations*, Stanford, CA: Stanford University Press, 1992, pp. 17-55.

Boyd, Gale, Stephen H. Karlson, Mark Neifer, and Marc Ross, "Energy Intensity Improvements in Steel Minimills," *Contemporary Policy Issues*, July 1993, 11(3): 88-100.

Boyd, Gale, Mark Neifer and Marc Ross, "Modeling Plant-Level Industrial Energy Demand with the Manufacturing Energy Consumption Survey (MECS) Database and the Longitudinal Research Database (LRD)" Argonne National Laboratory, ANL/EAIS/TM-96, Argonne, IL, 1992.

Caves, Richard E., *Industrial Efficiency in Six Nations*, Cambridge, MA: MIT Press, 1992.

Caves, Richard E. and D.R. Barton, *Efficiency in U.S. Manufacturing Industries*, Cambridge, MA: MIT Press, 1990.

DeCanio, Stephen J., "Barriers Within Firms to Encourage Energy-Efficient Investments," *Energy Policy* 21(9): 906-914, 1993.

Farina, Richard., *Been Down So Long It Looks Like Up to Me*, New York: Random House, 1966.

Ferrier, G.D. and J.G. Hirschberg, "Climate Control Efficiency," *The Energy Journal*, 1992, 13(1): 37-54.

Frantz, Roger S., "Ex-Ante and Ex-Post Criticisms of X-Efficiency Theory and Literature," in Weiermair, Klaus, and Mark Perlman, editors, *Studies in Economic Rationality: X-Efficiency Examined and Extolled*, Ann Arbor, MI: U. of Michigan Press, 1990, pp. 43-62.

Fried, Harold O., C.A. Knox Lovell, and Shelton S. Schmidt, editors, *The Measurement of Productive Efficiency: Techniques and Applications*, New York: Oxford University Press, 1993.

Hassett, Kevin A. and Gilbert E. Metcalf, "Energy Conservation Investment: Do Consumers Discount the Future Correctly?," *Energy Policy*, June 1993, 21(6): 710-716.

Hausman, Jerry (1979). "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables," *Bell Journal of Economics*, 10: 33-54.

Jaffe, Adam B. and Robert N. Stavins, "The Energy Efficiency Gap: What Does It Mean?," *Energy Policy*, 1994 forthcoming.

Kolbe, Larry and Amey Moot, "Demand-Side Management Program Evaluation," *Electric Power Research Institute Capital Budgeting Notebook*, Palo Alto, CA, 1992.

Leibenstein, Harvey, "Allocative Efficiency vs. 'X-Efficiency,'" *American Economic Review*, 1966, 56: 392-415.

Murrell, Peter, "Can Neoclassical Economics Underpin the Reform of Centrally Planned Economies," *Journal of Economic Perspectives*, Fall 1991, 5(4): 59-76.

Nichols, Albert L. "Demand-Side Management: Overcoming Market Barriers or Obscuring Real Costs?" *Energy Policy*, 1994 forthcoming.

Smith, Vernon L., "Economics in the Laboratory," *Journal of Economic Perspectives*, Winter 1994, 8(1): 113-131.

Sutherland, Ronald J. "Market Barriers to Energy-Efficient Investments," *The Energy Journal*, 1991, 12(3): 15-34.

Thaler, Richard H. *The Winner's Curse: Paradoxes and Anomalies of Economic Life*, New York: The Free Press (Macmillan), 1992.

Train, Kenneth. "Discount Rates in Consumers' Energy-Related Discussions: A Review of the Literature," *Energy* 10(12): 1243-1253, 1985.

ENDNOTES

1. With apologies to Richard Farina, Been Down So Long It Looks Like Up to Me, New York: Random House, 1966.

2. Energy Modeling Forum, Stanford University. Although the views expressed here are my own, the discussions of the Energy Modeling Forum working group on energy conservation helped to identify the key issues in the debate on the energy efficiency "gap." I am greatly indebted to the high quality of that group's discussion. With the usual disclaimers, I would also like to express my gratitude to Gale Boyd, Stephen Elliot, Blake Johnson, Alan Sanstad, and Ronald Sutherland for some helpful discussions of specific points.

3. The concept "technical inefficiency" might loosely be called consumer (producer) inertia but actually includes many factors, as explained below and should be clearly differentiated from "allocative inefficiency," which economists have studied extensively.

4. While many analysts emphasize energy price risks, other risks related to performance and reliability may be more important.

5. These costs apply even when consumers are risk neutral.

6. Risks on future returns would have to be greater in the short run than over longer horizons to explain this anomaly. Such might be the case for stocks, whose returns are much more volatile over the next year or two than over a 10 or 20-year horizon (Bogle 1994, pp. 30-31). However, many risks appear to increase over time, as with the long-term bondholder who must be rewarded with higher rates of return to compensate him for greater inflation risk.

7. This frontier is an isoquant from the industry's production function. A recent and useful explanation of the relevant techniques can be found in Fried et al (1993).

8. Inefficiency increased with tariff protection. It increased or decreased with the first two factors, depending upon the study.

9. I have adopted this example, with some modification, from a presentation by Kolbe and Moot (1992).