Energy Policy 1994 **22** (10) 804–810

The energy-efficiency gap

What does it mean?

Adam B Jaffe and Robert N Stavins

As renewed attention has been given by policy makers to energy conservation issues, it has frequently been asserted that an energy-efficiency gap exists between actual and optimal energy use. The critical question is how to define the optimal level of energy efficiency. This paper seeks to disentangle some confusing strands of argument that are frequently brought to bear on this question, by identifying the major conceptual issues that determine the set of feasible answers. We identify five separate and distinct notions of optimality: the economists' economic potential, the technologists' economic potential, hypothetical potential, the narrow social optimum and the true social optimum. Each of these has associated with it a corresponding definition of the energy-efficiency gap. Our analysis demonstrates that necessary preconditions for identifying the right measure of the energy-efficiency gap include understanding and disentangling market failure and non-market failure explanations for the gradual diffusion of energy-efficient technologies.

Keywords: Energy-efficiency gap; Market barriers; Market failures

Renewed attention is now being given by policy makers to issues of energy conservation, because of concerns about global climate change arising from the greenhouse effect, largely a function of fossil fuel combustion as a source of carbon dioxide (CO_2) emissions. In such policy contexts, it is frequently asserted that an energy-efficiency gap exists between current or expected future energy use, on the one hand, and optimal current or future energy use, on the other hand. Clearly, estimates of the magnitude of this gap will be determined mainly by how optimal behavior is defined. In the context of simulation analyses, the critical question becomes one of defining the optimal scenario.[1]

Adam B Jaffe is with the Department of Economics, Brandeis University, Waltham, MA 02254, USA; Robert N Stavins is with the John F Kennedy School of Government, Harvard University, Cambridge, MA 02138, USA.

Our purpose in this paper is to identify the major conceptual issues that surround different approaches to defining an optimal energy use scenario, and to explore the relationship between alternative notions of optimality and concepts that are widely used in discussions about the energy-efficiency gap. In this way, we implicitly consider the appropriate definition of the energy-efficiency gap itself.

The primary motivation for our investigation is to provide some guidance for public policy regarding energy and energy technologies. We adopt the standard economic approach of defining good public policy to be that which maximizes the appropriately weighted sum of the values of goods and services (including, of course, intangibles) enjoyed by society throughout time. Thus, we do not consider energy efficiency as a goal in itself, but only as a means to the end of overall efficient (and equitable) resource allocation.[2] To the extent that energy generation and/or use creates environmental externalities – one of the primary reasons why energy use is a public policy concern in the first place – we assume that such effects can be incorporated in the analysis by placing appropriate values on the environmental disamenities associated with energy generation and use.

The energy paradox

The crux of the debate surrounding the efficiency gap lies in differing interpretations of what has been called the paradox of gradual diffusion of apparently cost-effective energy-efficiency technologies.[3] Why are compact fluorescent light bulbs, improved thermal insulation materials and energy-efficient appliances not more widely used? For our analysis, we take it as given that such technologies (and processes) exist – ones that simple net present value calculations show to be cost-effective at current prices but which enjoy only limited market success.[4] We use the phrase market barriers to

0301-4215/94/10 0804-07

refer to any factors that may account for this apparent anomaly.[5] Differing views about the nature of these barriers lead to fundamentally different views about optimal energy use.

First of all, we should note that – paradox or no paradox – the diffusion of economically superior technologies is typically gradual.[6] Awareness of this empirical reality should make the existence of the so-called paradox much less perplexing, but it does not answer the question of whether the optimal rate of diffusion is greater than the observed rate.[7] As we will see below, at a theoretical level it is relatively easy to come up with reasons why technology diffusion will not, in general, occur at the socially optimal rate. But, if the purpose of measuring the efficiency gap is to identify desirable government policy interventions, then what we need to know is whether the market barriers that cause slow diffusion can be mitigated by government intervention in such a way that overall resource allocation is improved. In this context, the appropriate definition of the optimal level of efficiency is that which is consistent with efficient overall resource use, including efficient use of government resources.

Those market barriers that might justify a public policy intervention to overcome them we denote as market failures. By implication, there may be some market barriers that are not market failures; in such cases, the existence of the paradox does not, in and of itself, call for policy responses. On the other hand, there may be some market failures that do not relate to the paradox, but that might still provide justifications for policy interventions. In other words, even if the paradox is resolved in the sense that people's decisions are found to be consistent with the costs they face, there could be other reasons – particularly environmental externalities – why the resulting behavior would deviate from the social optimum. We consider that class of market failures below, but first we examine both market failure and non-market failure explanations of the paradox of gradual diffusion of energy-efficiency technologies.

Market failure explanations of the paradox

Several sources of potential market failure may affect energy conserving technology adoption rates.[8] Three of these relate to the availability of information. First, information has important public good attributes: once created it can be used by many people at little or no additional cost. It may be difficult or impossible for an individual or firm that invests in information creation to prevent others who do not pay for the information from using it. It is well known that such public goods will tend to be underprovided by ordinary market activity. Second, if the act of adopting a new technology is, itself, a source of useful information for others, then the act of adoption creates a positive externality by providing information to others for which the adopter is unlikely to be compensated. This (positive) externality is another form of market failure. Third, if the potential adopter is not the party that pays the energy bill, then good information in the hands of the potential adopter may not be sufficient for optimal diffusion; adoption will only occur if the adopter can recover the investment from the party that enjoys the energy savings.[9] Thus, if it is difficult for the possessor of information to convey it credibly to the party that benefits from reduced energy use, a principal/agent problem arises. This is another potential form of market failure. A home builder, acting as agent for the ultimate home buyer, may have incentives to take actions that are different from those the principal would prefer.[10]

Non-market failure explanations of the paradox

Non-market failure explanations of the energy paradox consist essentially of explaining why observed behavior is indeed optimal from the point of view of energy users.[11] To be useful, such explanations must advance beyond the tautological assertion that if the observed rate of diffusion is less than the calculated optimal rate, there must be some unobserved adoption costs that would modify our calculations of what is optimal.[12] One such explanation is that uncertainty about future energy prices and the actual savings from the use of energy technologies, combined with the irreversible nature of the efficiency investment, make the appropriate discount rate for analyzing the net present value of energy savings significantly greater than is typically used in the calculations that suggest the existence of a paradox.[13] Note that uncertainty, in contrast to imperfect information, is not a source of market failure in and of itself. It is reasonable and appropriate for individuals to take uncertainty into account in making investment decisions, and to apply relatively high discount rates to irreversible investments whose returns are uncertain.[14] To the extent that consumers' true discount rates were high for these reasons, this would not represent a market failure.

A second potential non-market failure explanation is provided by the possibility that qualitative attributes of new technologies may make them less desirable than existing, less efficient technologies. An obvious example is the difference in hue between fluorescent and incandescent lighting. Third, there are typically costs of adoption (including the non-public good part of information acquisition) that are not included in simple cost-effectiveness calculations. It is by no means costless to learn how a technological improvement fits into one's home or firm or to learn about reliable suppliers.[15] Even after basic information about a technology has been disseminated, the purchase price of a new product is only a lower bound on its adoption cost.

A third possible non-market failure explanation is associated with the fact that even if a given technology is cost-effective on average, it will mostly probably not be so for some individuals or firms. If the relevant population is heterogeneous with respect to the amount of energy it uses, for example, even a technology that looks very good for the average user will be unattractive for a portion of the population. Hence, we can also interpret cost-effectiveness calculations to mean that the technology is profitable for the mean household or firm. Depending on how the paradox is measured, such underlying heterogeneity can provide yet another non-market failure explanation.[16]

Finally, another way of characterizing the energy-efficiency gap itself is to say that inertia exists in consumers' adoption behavior. This is not an explanation, of course, but simply an additional characterization of the problem. The important question that remains is what causes the inertia. To the degree that it is caused by real costs of adoption borne by potential adopters, we are returned to one of the types of non-market failure described above. To the degree that the inertia is caused by some sort of (informational) market failure, we can then ask which of the types previously discussed is responsible.

Market failure issues that do not bear on the paradox

There are also market failure issues that do not relate to the paradox – that is, that do not help explain non-adoption at current prices – but which are relevant to policy debates about the energy-efficiency gap. First, actual energy prices, particularly for electricity, may differ from marginal social cost because of subsidies and because of pricing based on average rather than marginal cost. During the 1980s, it was widely perceived that the incremental costs of increasing electricity supplies were significantly greater than the average costs of existing electrical capacity. Since the prices consumers pay are typically based on historical average costs, it was frequently suggested that consumers faced inadequate incentives to conserve electricity. Each kilowatt of capacity that did not need to be built saved society much more than it saved the persons whose conservation decisions reduced the need for such new capacity. If true, this would be one reason why public policy should seek to promote greater energy efficiency than private individuals choose on their own (although it is not an explanation for the paradox, since the latter is a statement about individuals' behavior given the actual prices individuals face).

While this argument remains conceptually valid, the widespread excess capacity that now characterizes the electricity industry in the USA renders the point moot in practice. It is simply no longer true that the incremental cost of capacity is well above the price paid by consumers. Indeed, given the availability of wholesale bulk power at prices closer to variable cost than total cost, it could be argued that the social value of reducing electrical utility loads is actually less than the prices paid by most consumers. Furthermore, this situation is likely to prevail in some parts of the country for the next five to ten years, depending upon the growth of electricity demand. Thus, regulatory distortions in the utility industry no longer provide an argument for policy interventions to foster energy conservation.[17]

One other major market failure suggesting that energy efficiency may be below the socially desirable level is associated with the environmental consequences of energy generation and use.[18] While much controversy surrounds the magnitude of the value of the environmental damages associated with energy use,[19] the direction of the effect is unambiguous. Whether or not there is a paradox, consumers face incentives to use more energy than is socially desirable if they do not bear the full costs of the pollution their energy use fosters.[20]

The relevance of discount rates

Much discussion surrounding the energy-efficiency gap is couched in terms of arguments about the appropriate discount rate to use in evaluating energy savings. It is useful to distinguish between a number of separate and distinct questions regarding the appropriate discount rate. First, we can attempt to estimate the implicit discount rate that consumers appear to be using when they make energy-efficiency decisions. Second, we can speculate about the correct discount rate for consumers to use in making such decisions. Third, we can compare the implicit rate or the correct rate with other private discount rates in the economy, particularly the rates used by utilities to evaluate energy supply options. Fourth and finally, we can ask whether any of these rates is equal to the social discount rate that ought to be applied in evaluating future energy savings when making public policy decisions.

First of all, the observation that consumers have high implicit discount rates[21] when they make energy-efficiency decisions is actually neither more nor less than a restatement of the existence of the energy paradox. To estimate implicit discount rates, we examine decisions actually made and calculate the discount rate that makes those decisions privately optimal, given our estimates of the costs and future energy savings of the investments and given the assumption that there are no important market failures impeding the adoption of efficient technologies. The paradox can thus be restated as the observation that these implicit discount rates appear to be much higher than other interest rates in the economy. To observe that implicit discount rates are high, however, says nothing about the reason why people are making the decisions they are. One possibility is that people are applying normal discount rates in the context of significant market failures; another possibility is that people actually utilize high discount rates in evaluating future savings.[22] The truth is probably some combination of the two. Thus, high implicit discount rates, on their own, are neither a market failure nor an explanation of observed behavior.

Further, if we observe only the relevant investment decisions themselves, it is fundamentally impossible to determine whether the observed behavior results from market failures or from truly high discount rates.[23] To make that distinction, we would need either to observe something that distinguishes the market failure explanations from the non-market failure ones, such as whether people with better information are more likely to purchase more efficient models, or else to calculate from some basis other than the investment decisions themselves what an appropriate discount rate for these investments would be.[24]

Thus, to investigate the energy-efficiency gap by changing the discount rates people use in making investment decisions amounts to assuming the answer. If the outcome with lower discount rates is considered the optimal result, then it is implicitly assumed that all market barriers are indeed market failures. Conversely, if we postulate that the optimal result is the one in which consumers are assumed to discount at observed high implicit rates, then it is implicitly assumed that there are no market failures.

To make this a bit more concrete, suppose that through calculations we find that consumers' true discount rates for some set of energy-efficiency investments are approximately 20% – higher than mortgage interest rates, but in the range of rates on personal credit cards. What would this suggest for public policy? If mandatory efficiency standards were being considered, for example, then requiring efficiency any greater than what someone with a 20% discount rate would choose would have the effect of making consumers worse off.[25] Even if a utility demand-side management (DSM) program were being evaluated, the appropriate discount rate could well be 20%, despite the fact that the utility might evaluate supply-side investments at its cost of capital, more probably in the range of 8–10%. The reason is that the risk created by these irreversible investments with uncertain returns is a real social cost that must be borne by someone – either the utility and its stockholders or its customers.[26] Thus, to the extent that high implicit discount rates correspond to truly high discount rates, rather than to market failures, there is nothing particularly wrong with those high rates, and they do not correspond to any efficiency gap that ought to be addressed by public policy.[27] Instead, we are returned to the question of whether relevant market failures exist.

Finally, there is the issue of the social discount rate. Some 30 years of academic and policy debates have failed to resolve the question of what the social rate of discount should be, even as a conceptual matter.[28] In any event, however, there is a strong argument that policies to increase energy efficiency can only be justified on the basis of the appropriate true private rate of discount, not some (lower) social rate. If the only reason why investment in energy-efficient technology is suboptimal is the divergence between the private and social discount rates, then it is certainly true that a theoretical argument could be made to support public policies to increase that investment. The problem is that this same argument would also apply with equal force to all other forms of investment – plant and equipment, research, education etc. It may be that the government should be doing more to encourage all these forms of investment, but that is surely an issue beyond the scope of energy policy.

Further, if the social rate of discount is really very low, such as 3%, for example, then the set of available investment opportunities that should be undertaken probably exceeds the nation's annual gross national product (GNP). Obviously, we cannot undertake all of these projects.[29] Given this, a policy prescription to increase investment in energy efficiency should be based on a conclusion that the (social) rate of return to this form of investment is higher than the rate of return available on other forms of investment in the economy. Otherwise, we could well be increasing this form of investment at the expense of others that are even more beneficial. Of course, saying that the social rate of return on energy-efficiency investments exceeds the social rate of return on other investments means precisely that current investment in energy efficiency is inadequate when evaluated at the appropriate private rate of discount.

Synthesis

We have attempted to disentangle the confusing and sometimes confused strands of argument that are typically brought to bear on discussions of the energy-efficiency gap. We can synthesize much of what we have said by examining graphically the concepts behind the major, alternative notions of the efficiency gap. In Figure 1 each plateau represents another version of an optimal scenario. Thus, different notions of the efficiency gap correspond respectively to distances

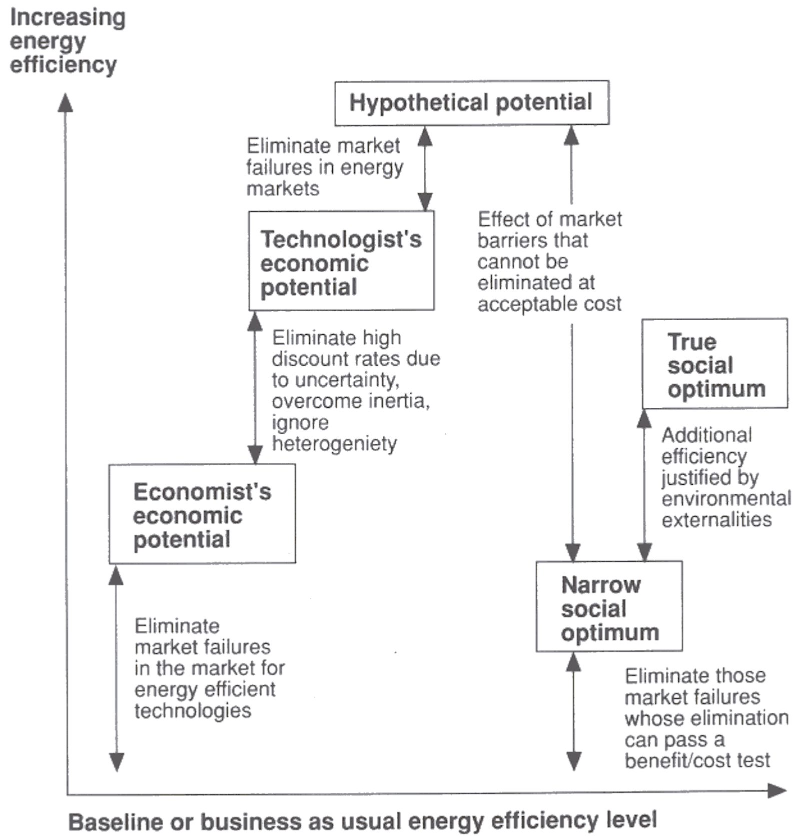

Figure 1. Energy-efficiency gaps.

between each of the plateaux and the horizontal axis – the base case or business as usual. The text in each column of the figure between successive plateaux describes the differences between them, and the height of each plateau provides a rough qualitative estimate of relative levels of efficiency (although it is theoretically possible that the ranking could be different from what is shown).

Our analysis suggests two distinct notions of economic potential and two distinct notions of social optimum. First of all, we use the general phrase economic potential to describe the degree of energy efficiency that would be achieved if various economic barriers were removed. Further, we describe the scenario in which we eliminate market failures in the energy technology market (such as inadequate provision of information) as the economists' economic potential. If we were also to remove non-market failure market barriers (such as high discount rates caused by uncertainty about payback), we would achieve the technologists' economic potential.[30] Clearly, if we do not distinguish between barriers to adoption that are market failures and those that are not, and choose to eliminate all barriers, we will achieve a higher estimate of potential efficiency.

Both notions of economic potential consider only issues relating to the energy paradox, that is, market barriers in the markets for energy technologies and processes. We also specify a hypothetical potential, in which the focus is broadened to include energy markets themselves, and any additional efficiency that might result from getting energy prices right is included. In this figure we show this plateau as only slightly higher than the previous one, corresponding to the view, explained above, that the incremental costs of energy supply do not currently differ greatly from consumers' costs in most parts of the country.

This hypothetical potential could be reached only if actual government programs were designed to overcome all market barriers. This is surely impossible, and probably not even desirable, since removal of some of these barriers would require government programs that were themselves excessively costly to implement. Hence, to find the narrow social optimum, we reduce energy efficiency by reintroducing those barriers whose cost of removal would be greater than the benefits that would thereby be created.[31] Alternatively, the narrow social optimum can be defined with reference to the baseline: it represents the energy efficiency achieved by instituting all available programs to encourage energy efficiency that pass an appropriate cost–benefit test.[32] As drawn, the figure represents a fairly pessimistic view of government's ability to institute efficient policies, because the narrow policy optimum is shown as being considerably lower than the economist's economic potential. If government's abilities in this regard are greater, then the narrow policy optimum ought to be correspondingly higher.

In the final column of the figure, we represent the true social optimum, which incorporates the additional energy conservation called for by internalizing the environmental effects of energy generation and use. As drawn, this is a large effect, corresponding to the apparent revealed preference of the policy process, which seems to give large weight to environmental considerations in energy policy decisions. Readers with different judgements can adjust this distance accordingly.[33]

How these alternative definitions of optimal energy efficiency and the energy-efficiency gap relate to specific empirical estimates depends, of course, on the assumptions that underlie those estimates. Many published studies of the potential for energy efficiency correspond to what we have labelled technologists' economic potential. That is, they assume that the resolution of the energy paradox must be that the simplest calculations are correct and that a host of market failures explain observed behavior. Certainly, estimates of efficiency potential derived from forcing low discount rates into analyses of consumers' and firms' decisions do not correspond to any policy relevant notion of optimal. Unfortunately, it is easier to explain what is wrong with existing approaches than to specify what is the right one. It is clear, however, that in order to understand what is wrong with some of the existing approaches to estimating the energy-efficiency gap and to begin the process of searching for the right measure, we first need to disentangle market failure and non-market failure explanations for observed decisions regarding energy-efficiency investments.

We gratefully acknowledge the comments on a previous version of this paper by Hill Huntington, Ronald Sutherland, and other colleagues at the Energy Modeling Forum. The authors alone, however, are responsible for any errors.

1When the energy-efficiency gap is between optimal future use and expected future use under some base case or business as usual scenario, as in the research of Working Group 13 (Energy Conservation) of the Energy Modeling Forum (EMF), there are also important questions associated with the appropriate definition of the base case. These issues, however, are beyond the scope of the present paper.

2For an examination of the contrast between the economist's implicit policy objective of efficient use of energy resources and the technologist's implicit policy objective of energy efficiency, see R J Sutherland, 'Energy efficiency or the efficient use of energy resources?', *Energy Sources*, Vol 16, 1994, pp 261–272.

3A Shama, 'Energy conservation in US buildings: solving the high potential/low adoption paradox from a behavioral perspective', *Energy Policy*, Vol 11, No 2, 1983, pp 148–167.

4For further discussion of these technologies and processes, see V Norberg-Bohm, *Potential for Carbon Dioxide Emissions Reductions in Buildings*, Energy and Environmental Policy Center, Global Environmental Policy Project Discussion Paper, John F Kennedy School of Government, Harvard University, 1990.

5Others have used the phrase market barriers even more broadly to include, for example, low energy prices that are a disincentive to the adoption of less energy-intensive technologies. See E Hirst and M Brown, 'Closing the efficiency gap: barriers to the efficient use of energy', *Resources, Conservation and Recycling*, Vol 3, No 4, 1990, pp 267–281.

6From the mechanical reaper of the nineteenth century (David, 1966), through hybrid corn seed (Griliches, 1957), chemical process innovations (Davies, 1979), steel furnaces (Oster, 1982), and optical scanners (Levin, Levin, and Meisel, 1987) in the 20th century, research has consistently shown that diffusion of new, economically superior technologies is never instantaneous, but typically follows instead an s-shaped or sigmoid curve, so that the adoption rate is initially slow, then faster, and then slower again as saturation is approached. See, respectively, P A David, 'The mechanization of reaping in the Ante-Bellum midwest', in H Rosovsky, ed, *Industrialization in Two Systems*, Harvard University Press, Cambridge, 1966, pp 3–39; Z Griliches, 'Hybrid corn: an exploration in the economics of technological change', *Econometrica*, Vol 25, No 4, 1957, pp 501–522; S W Davies, 'Inter-firm diffusion of process innovations', *European Economic Review*, Vol 12, No 4, 1979, pp 299–317; S Oster, 'The diffusion of innovation among steel firms: the basic oxygen furnace', *Bell Journal of Economics*, Vol 13, No 1, 1982, pp 45–56; and S G Levin, S L Levin and J B Meisel, 'A dynamic analysis of the adoption of a new technology: the case of optical scanners', *Review of Economics and Statistics*, Vol 69, No 1, 1987, pp 12–17.

7To some degree, responses to the observation that new technologies always diffuse slowly may depend on ideological perspectives. Liberals may be inclined to note that if technology diffusion is typically not optimal, then there is no basis for presuming that the market is working efficiently, and hence no reason to eschew government intervention designed to improve it. Conservatives, on the other hand, may be more inclined to maintain that the government cannot possibly worry about trying to get the diffusion rate right for all technologies, and that the government should hence stay out altogether, even if the *laissez-faire* result is not optimal.

8For a more comprehensive and formal analysis of alternative explanations of the energy paradox, see A B Jaffe and R N Stavins, 'The energy paradox and the diffusion of conservation technology', *Resource and Energy Economics*, Vol 16, No 2, 1994, pp 91–122, 1994.

9For example, in construction of new homes, builders may have difficulty conveying the benefits of energy conserving technologies to prospective buyers because these technologies (and their future energy use consequences) are not observable. Likewise, landlords may not be able to recover all of the value of such investments (in the form of higher rents) where renters pay fuel bills; and tenants who make these investments in cases where the landlord pays the energy bill may not be able to get reduced rents.

10Some analysts have identified the transaction costs associated with adopting new technologies (such as costs associated with searching and learning) as an additional category of market failure. See M D Levine, E Hirst, J G Koomey, J E McMahon and A H Sanstad, *Energy Efficiency, Market Failures, and Government Policy*, Lawrence Berkeley Laboratory, University of Califoria, Working Paper, March 1994. Other than with regard to the public good dimensions of learning, such adoption costs are more aptly described as part of the typically unmeasured costs of technology and hence, if anything, constitute a category of non-market failure explanations of the paradox, as discussed below.

11See, for example, P L Joskow and D B Marron, 'What does a negawatt really cost? Evidence from utility conservation programs', *Energy Journal*, Vol 13, No 4, 1992, pp 41–74; and A Nichols, *Estimating the Net Benefits of DSM Programs Based on Limited Information*, National Economic Research Associates, Cambridge, MA, 1992.

[12]A slightly less tautological approach would be to maintain the assumption – based on general experience – that private agents act in their own interests, unless it can be shown that specific market failures exist. Under this approach, it is quickly concluded that no paradox exists, unless and until it can be demonstrated that the theoretical market failures discussed above are important in explaining observed behavior. This amounts to a policy prescription that markets should be considered innocent until proven guilty. We find no real basis for such a prescription. On the other hand, it is important to keep in mind that although market failure may be a necessary condition for government policy intervention, it is not a sufficient condition. For government policy to be desirable, it must be the case that a market failure exists and that there exists a government policy that, in overcoming the market failure, generates benefits in excess of the costs of implementation.

[13]R J Sutherland, 'Market barriers to energy-efficiency investments', *The Energy Journal*, Vol 12, No 3, 1991, pp 15–34.

[14]K A Hassett and G E Metcalf, *Energy Tax Credits and Residential Conservation Investment*, National Bureau of Economic Research, Cambridge, MA, Working Paper No 4020, March 1992. For broader treatments of this issue, see R S Pindyck, 'Irreversibility, uncertainty, and investment', *Journal of Economic Literature*, Vol 29, No 3, 1991, pp 1110–1152; and A Dixit, 'Investment and hysteresis', *Journal of Economic Perspectives*, Vol 6, No 1, 1992, pp 107–132.

[15]Some have argued that not only costly information acquisition but also biased estimates of likely energy savings play a role. Consumers may not believe experts' assessments of the benefits of new technologies. On the other hand, the bias may go in the opposite direction, since some studies indicate that consumers systematically overestimate energy savings associated with some types of new technologies. See P C Stern, 'Blind spots in policy analysis: what economics doesn't say about energy use', *Journal of Policy Analysis and Management*, Vol 5, No 2, 1986, pp 200–227.

[16]A B Jaffe and R N Stavins, 'Evaluating the relative effectiveness of economic incentives and direct regulation for environmental protection: impacts on the diffusion of technology', paper presented at the National Bureau of Economic Research Summer Institute Workshop on Environmental Economics, August 1991.

[17]With respect to technology incorporated in new buildings, it can be argued that these buildings are likely to last long into the future, at which time new electrical capacity will once again be needed, so that consumers choosing energy-inefficient technology are imposing a social cost above their own energy costs, albeit one that is remote in time. With respect to retrofit investments, however, even this argument is unavailable, since these investments can simply be postponed until such time as they are needed to forestall the need to expand capacity.

[18]It has also been suggested that there are externalities associated with the economic and military security costs resulting from domestic US dependence on imported oil from politically unstable regions. For this argument to be valid, it would have to be the case that, at the margin, these national security costs are reduced if oil consumption is marginally reduced. This seems unlikely to be the case. See D R Bohi and M A Toman, 'Energy security: externalities and policies', *Energy Policy*, Vol 21, 1993, pp 1093–1109.

[19]See, for example, W K Viscusi, W A Magat, A Carlin and M K Dreyfus, 'Environmentally responsible energy pricing', *Energy Journal*, forthcoming, 1994.

[20]As we discuss below, this may not be true if the government is already intervening to change the energy efficiency of products available in the market place (or is otherwise intervening to control the pollution associated with energy use). For example, Corporate Average Fuel Economy (CAFE) standards result in the production and sale of cars that are more efficient than consumers would otherwise demand. Depending on the actual magnitude of the environmental externality, it is possible that the efficiency of automobiles is greater than would be socially optimal.

[21]There is substantial empirical evidence that consumers do use high implicit discount rates in making energy conservation investment decisions. Hausman (1979) estimated that consumers used average implicit discount rates of 20% for purchasing room air conditioners with substantial variation by income class; and Dubin and McFadden (1984) found average implicit discount rates of 20% for space heating and water heating investments, again with significant variation by income. In a comment on Hausman, Gately (1980) estimated discount rates of 45 to 300% for refrigerators. Likewise, Ruderman *et al* (1987) found personal implicit discount rates as low as 20% and as high as 800% for heating and cooling equipment and for residential appliances. See, respectively, J A Hausman, 'Individual discount rates and the purchase and utilization of energy-using durables', *The Bell Journal of Economics*, Vol 10, No 1, 1979, pp 33–54; J A Dubin and D L McFadden, 'An econometric analysis of residential electric appliance holdings and consumption', *Econometrica*, Vol 52, No 2, 1984, pp 345–362; G Dermot, 'Individual discount rates and the purchase and utilization of energy-using durables: comment', *Bell Journal of Economics*, Vol 10, No 1, 1980, pp 373; and H Ruderman, M D Levine and J E McMahon, 'The behavior of the market for energy efficiency in residential appliances including heating and cooling equipment', *The Energy Journal*, Vol 8, No 1, 1987, pp 101–124.

[22]When we say actually, we do not mean literally. It is not necessary that people know how to perform net present value calculations; rather, they make decisions balancing costs today and costs tomorrow, with those trade offs incorporating their true discount rate.

[23]For an example of this in the context of an empirical analysis, see A B Jaffe and R N Stavins, 'Dynamic incentives of environmental regulations: the effects of alternative policy instruments on technology diffusion', *Journal of Environmental Economics and Management*, forthcoming, 1995.

[24]For an attempt at this latter approach, see, *op cit*, Ref 14, Hassett and Metcalf.

[25]*Op cit*, Ref 13.

[26]It is possible, of course, that particular supply-side investments are also very risky. The utility cost of capital does, however, capture the average riskiness of the asset base held by the firm.

[27]A more complicated possibility is that conservation investments have an inherent riskiness comparable with supply-side investments, but that consumers are less able to mitigate such risks through pooling across projects or other forms of diversification. If so, then the observed (higher) discount rate would be appropriate for evaluating standards that force consumers to make the investments themselves, but the (lower) utility cost of capital would be appropriate for evaluating utility investments.

[28]For an overview of the issues, see R C Lind, ed, *Discounting for Time and Risk in Energy Policy*, Resources for the Future, Washington, DC, 1982.

[29]If we take seriously the notion that public policy should have as its objective undertaking all investments whose rates of return exceed the social discount rate, then the reasonable prescription is to increase investment, beginning with the investments with the highest rates of return. Increasing investment reduces consumption; long before investment opportunities with rates of return above 3% were exhausted, consumption would be so significantly reduced that the social discount rate would rise. We can imagine increasing investment up to the point where the private and social rates of return coincide, but it is impossible to know how much greater investment would be at that point or what the discount rate would be.

[30]In the figure, we treat inertia as a non-market failure. Hence, it represents one part of what differentiates the technologists' notion of economic potential from the economists' notion.

[31]The degree of energy efficiency in the technologists' or economists' economic potential must be greater than or equal to the narrow socially optimal level, since overcoming some of the market barriers may be so costly that it is not worth doing. Note that the costs of eliminating a market barrier include government expenditures, regulatory burdens, loss of utility from degraded product performance etc.

[32]By definition, for non-market-failure market barriers, it will not be possible to design countermeasures that pass a cost–benefit test. If there is no market failure removing that barrier creates only negative benefits.

[33]This discussion ignores the reality that the baseline efficiency already reflects policy measures motivated by environmental concerns. As emphasized previously, it is possible that these measures already achieve as much or more additional efficiency as would be justified by the environmental externalities of energy use. If this is the case, the social optimum would be no greater than the narrow social optimum.