Economics of Energy Efficiency

ADAM B. JAFFE
Brandeis University and National Bureau of Economic Research
Waltham, Massachusetts, United States
RICHARD G. NEWELL
Resources for the Future
Washington, D.C., United States
ROBERT N. STAVINS
Harvard University
Cambridge, Massachusetts, United States

Glossary

externality An economically significant effect of an activity, the consequences of which are borne (at least in part) by parties other than the party who engages in the activity, and which are not accounted for through trade.

diffusion The gradual adoption of new process or product innovations by firms and individuals.

innovation The initial market introduction or commercialization of new process or product inventions.

invention The development and creation of a prototype new idea, process, or piece of equipment.

market barriers Disincentives to the diffusion and/or use of a good, such as high costs or prices, which may or may not represent market failures.

market failures The failure of private markets to provide certain goods at all or at the most desirable level, typically arising from a situation in which multiple parties benefit from a good without decreasing one another's benefits, and in which those who have paid for the good cannot prevent others from benefiting from it.

technological change The process of invention, innovation, and diffusion whereby greater and/or higher quality outputs can be produced using fewer inputs.

This article reviews economic concepts relevant to decision making about energy efficiency investments and related policy choices. We describe economic perspectives on the process of energy-saving technological change, the distinction between market failures and market barriers in energy-using product markets, and the important role that discounting plays in this area.

1. OVERVIEW

Energy efficiency is defined here to mean energy services provided per unit of energy input (for example, gallons of water heated to a specified temperature per British thermal units of natural gas input). Within this framework, energy efficiency is conceived primarily at the disaggregated, product level, rather than at a more aggregated sectoral level. As with virtually all economic problems, the economics of energy efficiency is at its heart a question of balancing of costs and benefits. For the individual energy user, this involves weighing the higher initial cost of purchasing energy-efficient products against the expected benefits of future cost savings when operating the products, among other considerations. Due to this difference in the timing of energy efficiency costs and benefits, issues related to discounting feature prominently in analyses of energy-efficient technology adoption. For suppliers of energy-using products, decisions regarding energy-efficient innovations likewise depend on the expected profits from such technology development. Profits from innovation depend in turn on the expected demand for energy-efficient technologies and the degree to which firms can appropriate the value created by their innovations.

On both the demand and supply side of this market for energy-efficient technology, potential market imperfections can lead to underinvestment in energy efficiency. This raises the possibility that corrective government policies could provide economic gains if they are practicable and provide net benefits after inclusion of all public and private implementation costs. The degree to which such opportunities exist in practice is the subject of significant debate. Finally, environmental pollution associated with energy production—particularly carbon dioxide emissions from fossil fuels—represents an additional reason why private markets might underprovide energy efficiency if energy users do not face the cost of any resultant environmental harm.

2. THE PROCESS OF TECHNOLOGICAL CHANGE

Many readers may be unfamiliar with the way economists typically view the process of technological change, thus it is useful to first establish this common understanding. Furthermore, to understand the potential for public policy to affect energy efficiency, it is also necessary to understand the process through which technology evolves: invention, innovation, diffusion, and product use. Invention involves the development of a new idea, process, or piece of equipment. This activity typically takes place inside the laboratory. The second step is the process of technology innovation, in which new processes or products are brought to market. Another way of describing this stage is commercialization. The third step is diffusion, the gradual adoption of new processes or products by firms and individuals, which then also decide how intensively to use new products or processes.

Tying this all together, it would be possible, for example, to think of a fundamentally new kind of automobile engine being invented. This might be an alternative to the internal combustion engine, such as a system dependent on energy-efficient fuel cells. The innovation step would be the work carried out by automobile manufacturers or others to commercialize this new engine—that is, bring it to market, then offer it for sale. The diffusion process, then, would reflect the purchase by firms and individuals of automobiles with this new engine. Finally, the degree of use of these new automobiles will be of great significance to demand for particular types of energy. The reason it is important to distinguish carefully among these different conceptual steps—invention, innovation, diffusion, and use—is that public policies can be designed to affect various stages and will have specific and differential effects.

2.1 Technology Adoption and Diffusion

Beginning at the end of the technological change process, research has consistently shown that diffusion of new, economically superior technologies is never instantaneous. The S-shaped diffusion path shown in Fig. 1 has typically been used to describe the progress of new technologies making their way into the marketplace. The figure portrays how a new technology is adopted at first gradually and then with increasing rapidity, until at some point its saturation in the economy is reached. The explanation for this typical path of diffusion that has most relevance for energy conservation investments is related to differences in the characteristics of adopters and potential adopters. This includes differences in the type and vintage of their existing equipment, other elements of the cost structure (such as access to and cost of labor, material, and energy), and their access to technical information. Such heterogeneity leads to differences in the expected returns to adoption and, as a result, only potential adopters for whom it is especially profitable will adopt at first. Over time, however, more and more will find it profitable as the cost of the technology decreases, quality improves, information about the technology becomes more widely available, and existing equipment stocks depreciate. The longevity of much energy-using equipment reinforces the importance of taking a longer term view toward energy efficiency improvements—on the order of decades (see Table I).

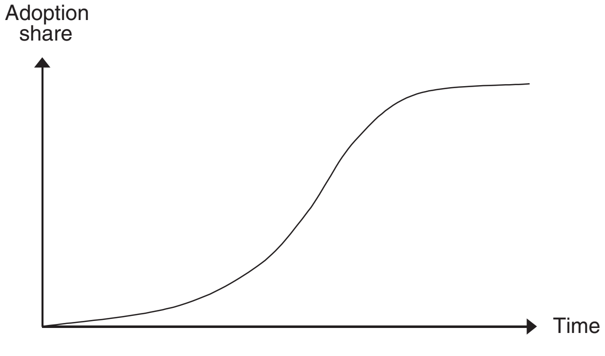

FIGURE 1 The gradual S-shaped path of technology diffusion.

Several studies have explored the effect of energy prices and technology adoption costs on energy efficiency investments. From a policy perspective, the effect of higher energy prices can be interpreted as suggesting what the likely effects of taxes on energy use (or carbon dioxide) would be and the effects of changes in adoption costs can be interpreted as indicating what the effects of technology adoption subsidies would be. As suggested by the economic cost–benefit paradigm, these studies have found that higher energy prices increase and adoption costs decrease the extent of adoption of energy-efficient technology (for example, building insulation, more efficient home appliances, or more efficient industrial motors). An additional interesting finding in this line of research is that the adoption of these technologies is more sensitive to the cost of the equipment than it is to the expected cost of energy. This implies that a policy of subsidizing the purchase of new efficient equipment may be more effective than a policy of taxing resource use, for policies that should in theory create the same magnitude of economic incentive.

There are at least three possible explanations for this divergence. One possibility is a behavioral bias that causes purchasers to focus more on up-front cost than they do on the lifetime operating costs of an investment. An alternative view is that purchasers focus equally on both, but uncertainty about future energy prices or whether they will face these costs (because they could move, for example) makes them give less weight to energy prices than they do to capital cost, which is known. A final interpretation might be that consumers have reasonably accurate expectations about future energy prices, and their decisions reflect those expectations, but the proxies for these expectations that are used by researchers are flawed, causing their measured effect to be smaller than their true effect. For example, studies often use current realized energy prices as a proxy for expected future energy prices. Current prices fluctuate more than expected future prices, however, leading to a downward bias in the coefficient on the energy price proxy relative to the true relationship with expected prices.

Although empirical evidence indicates that subsidies may be more effective than comparable taxes in encouraging technology diffusion, it is important to recognize some disadvantages of such subsidy approaches. First, unlike energy prices, adoption subsidies do not provide incentives to reduce utilization. Second, technology subsidies and tax credits can require large public expenditures per unit of effect, because consumers who would have purchased the product even in the absence of the subsidy still receive it.

TABLE I

Technology Diffusion and the Rate of Capital Stock Turnover[a]

Type of asset	Typical service life (years)
Household appliances	8–12
Automobiles	10–20
Industrial equipment/machinery	10–70
Aircraft	30–40
Electricity generators	50–70
Commercial/industrial buildings	40–80
Residential buildings	60–100

[a] Technology diffusion is closely related to the concept of "capital stock turnover," which describes the rate at which old equipment is replaced and augmented by new. New equipment can be purchased either to replace worn out and obsolete units or as a first-time purchase. A primary driver of replacement purchases for durable energy-using products is the useful lifetime of the product. The rate of economic growth is also important, especially for first-time durable-goods purchases; the rate of home construction is particularly relevant for residential equipment. The typical lifetimes for a range of energy-using assets illustrate that the appropriate time frame for thinking about the diffusion of many energy-intensive goods is on the order of decades.

2.2 Technology Innovation and Invention

Now it is possible to move back in the process of technological change from diffusion to innovation. In the energy efficiency area, it is helpful to think of the innovation process as affecting improvements in the attributes or characteristics of products. In Fig. 2, this process is represented as the shifting inward over time of a curve representing the trade-offs between different product characteristics for the range of products available on the market. On one axis is the cost of the product, and on the other axis is the energy flow (use; that is, the energy intensity) associated with a product. The downward slope of the curves indicates the trade-off between equipment cost of energy efficiency. Innovation means an inward shift of the curve—greater energy efficiency at the same cost, or lower cost for a given energy efficiency.

As with technology diffusion, studies have shown that increases in the price of energy have induced technological innovations in energy efficiency of commercialized products, such as household appliances (e.g., air conditioners, water heaters), automobiles, tractors, and jet aircraft. Moving back even further in the process of technological change to

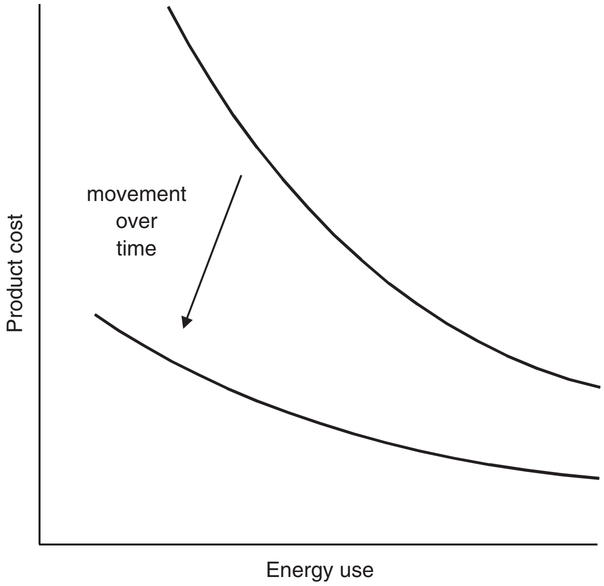

FIGURE 2 Innovation in product characteristics.

examine invention, other studies have analyzed U.S. patent data, finding that the rate of energy-conserving device patent applications (e.g., for waste heat, heat pumps, solar energy, and fuel cells) is significantly and positively associated with the price of energy.

3. UNDERSTANDING THE ENERGY EFFICIENCY GAP

Renewed attention is now being given by policymakers to energy efficiency, primarily due to concerns about global climate change associated with carbon dioxide emissions from fossil fuels. Much attention has been placed on the role that technological improvements in energy efficiency can play in reducing carbon emissions and in lowering the cost of those reductions. In such policy contexts, it is frequently asserted that there exists an "energy efficiency gap" between current energy use, on the one hand, and optimal energy use, on the other hand. From this perspective, the energy efficiency gap discussed earlier can now be thought of as a debate mainly about the gradual diffusion of energy-saving technologies that seem to be cost-effective.

Clearly, an estimate of the magnitude of this gap will be determined mainly by how optimal behavior is defined. Several major conceptual issues surround different approaches to defining an optimal energy-use scenario, with implications for guidance regarding public policy for energy and energy technologies. The standard economic approach defines "good" public policy to be that which maximizes the appropriately weighted sum of the values of goods and services enjoyed by society throughout time (including intangibles such as the environment). From this perspective, energy efficiency is not a goal in itself, but only a means to the end of overall economically efficient (and equitable) resource allocation. To the extent that energy generation and/or use creates environmental problems—one of the primary reasons why energy use is a public policy concern in the first place—such effects can in principle be incorporated in analyses by placing appropriate values on the environmental and nonenvironmental benefits and costs associated with energy generation and use.

The crux of the debate surrounding the energy efficiency gap lies in differing interpretations of what has been called the paradox of gradual diffusion of apparently cost-effective energy-efficient technologies. Why are compact fluorescent light bulbs, improved thermal insulation materials, and energy-efficient appliances not more widely used? Many studies have demonstrated that there exist such technologies (and processes)—ones that simple net-present-value calculations show to be cost-effective at current prices and market interest rates, but which enjoy only limited market success. The phrase market barriers has been used to refer to any factors that may account for this apparent anomaly. Others have used the phrase market barriers even more broadly to include, for example, low energy prices that are a disincentive to the adoption of more energy-efficient technologies. Differing views about the nature of these barriers lead to fundamentally different views about optimal energy use and the role of government policy.

As already noted, the diffusion of economically superior technologies is typically gradual. Awareness of this empirical reality should make the existence of the energy efficiency gap much less perplexing, but it does not answer the question of whether the optimal rate of diffusion is greater than the observed rate. To some degree, responses to the observation that new technologies always diffuse slowly depend on subjective perspectives. Some are 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. On the other hand, others are more inclined to maintain that the government cannot possibly worry about trying to optimize the diffusion rate for all technologies, and that the government should hence stay out altogether, even if the no-intervention result is not optimal.

As will be described further here, at a theoretical level there are several reasons why technology diffusion will not, in general, occur at an economically efficient rate. But, if the purpose of assessing the efficiency gap is to identify desirable government policy interventions, then it is necessary to know whether the market barriers that cause slow diffusion can be mitigated by government intervention such that overall resource allocation is improved. In this context, the appropriate definition of the optimal level of energy 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 are referred to in economics parlance as market failures. Normative economics teaches that in the presence of market failures, unfettered markets may not operate to produce outcomes that are socially optimal. When analysts speak of no-cost climate policies based on energy efficiency enhancement, they are often implicitly or explicitly assuming the presence of market failures in energy efficiency. By implication, there are some market barriers—such as "low" energy prices, "high" technology costs, and "high" discount rates—that are not necessarily market failures. In such cases, the existence of an energy efficiency gap does not, in and of itself, call for policy responses. On the other hand, there are also some market failures that do not relate to problems in individual decision-making per se, but that might still provide justification for policy interventions. In other words, even if the energy efficiency 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 pollution from energy use—why the resulting behavior would deviate from the social optimum. These classes of market failures and non-market-failure explanations of the energy efficiency gap are considered next.

4. MARKET FAILURES

In this section, the focus is on several sources of potential market failure that may affect energy-conserving technology adoption rates. Three of these relate to the availability of information. There are also market failure issues that do not help explain nonadoption at current technology costs and energy prices, but which are still relevant to policy debates about the energy efficiency gap. These include environmental externalities, broader innovation and adoption externalities, and issues related to energy supply pricing and national security.

4.1 Inadequate 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 inherently 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. It takes time for the many potential users to learn of the new technology, try it, adapt it to their circumstances, and become convinced of its superiority. An important mechanism in this learning process is the observation of the adoption of the new technology by others. If a neighbor or competitor tries something new, and others see that it works, it becomes much safer and easier for the observers to try it. Hence the adopter of a new technology creates a positive externality for others, in the form of the generation of information about the existence, characteristics, and successfulness of the new technology. This phenomenon is often called learning by using. This (positive) externality is another form of market failure.

Incomplete information can also foster principal-agent problems, as when a builder or landlord chooses the level of investment in energy efficiency in a building, but the energy bills are paid by a later purchaser or a tenant. If the purchaser has incomplete information about the magnitude of the resulting energy savings, the builder or landlord may not be able to recover the cost of such investments, and hence might not undertake them. 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 what the principal would prefer. The extent to which this market failure is significant in practice is not well established.

4.2 Environmental Externalities

Economic analysis of environmental policy is based on the idea that the potentially harmful consequences of economic activities on the environment constitute an externality. An externality is an economically significant effect of an activity, the consequences of which are borne (at least in part) by a party or parties other than the party who controls the externality-producing activity. An electricity generator that pollutes the air imposes a cost on society. The firm that owns the generator has an economic incentive to use only as much fuel as it can productively employ, because this input is costly to the firm. In jargon, the cost to society of having some of its fuel used up is internalized by the firm, because it has to pay for that input. But the firm does not (in the absence of appropriate environmental policy intervention) have an economic incentive to minimize the external costs of pollution associated with that fuel.

Although the details and refinements are important, all environmental policies, at their core, are designed to deal with this externality problem, either by internalizing environmental costs so that polluters will make efficient decisions regarding their consumption of environmental inputs, or else by imposing from the outside a level of environmental pollution that policymakers believe to be more efficient than that otherwise chosen by firms. If environmental externalities are not fully addressed by environmental policy—for political or other reasons—the resulting level of energy efficiency will likely be too low. But 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 environmental externalities, it is possible that the efficiency of automobiles is greater than would be socially optimal.

4.3 Broader Innovation and Adoption Externalities

In addition to the externality associated with pollution, innovation and diffusion are both characterized by externalities, as well as other market failures. In the case of pollution, the problem is that a polluter imposes costs on others, and hence has an inadequate incentive, from a social perspective, to reduce those costs. With respect to technology, the problem is the reverse. A firm that develops or implements a new technology typically creates benefits for others, and hence has an inadequate incentive to increase those benefits by investing in technology. Pollution is a negative externality, and so the invisible hand allows too much of it. Technology creates positive externalities, and so the invisible hand produces too little of it.

With respect to innovation, the positive externality derives from the public-good nature of new knowledge. Whereas patents and other institutions try to protect firms' investments in innovation, such protection is inherently imperfect. A successful innovator will capture some rewards, but those rewards will always be only a fraction—and sometimes a very small fraction—of the overall benefits to society of the innovation. Hence innovation creates positive externalities in the form of knowledge spillovers for other firms, and spillovers of value for the users of the new technology.

In addition, production costs tend to fall as manufacturers gain production experience in a process commonly called learning by doing. If this learning spills over to benefit other manufacturers, it can represent an additional externality. In any event, the existence of these innovation and adoption externalities suggests a rationale for public support of research and development in general, not for energy-efficient technology in particular. An argument particularly supporting energy-efficient technologies in this regard would require stating that spillovers related to these technologies were somehow greater than for technologies more generally in the economy.

4.4 Market Failures in Energy Supply

4.4.1 Average-Cost Pricing

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. Because 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 this argument remains conceptually valid in part, several things have changed since the 1980s that weaken the argument considerably. First, the deregulation of electricity production for about 40% of the U.S. population has meant a movement toward competitive electricity prices; this trend will continue. Second, there is widespread excess capacity in the electricity industry in most of the United States, although there is substantial regional variation. Thus, it is simply no longer true that the incremental cost of capacity is well above the price paid by most consumers. Indeed, given the availability of wholesale bulk power at prices closer to variable cost than to 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 for many years, depending on the growth of electricity demand. Thus, regulatory distortions in the utility industry no longer provide an argument for policy interventions to foster energy conservation. One caveat, however, is that at peak demand periods, the marginal cost of supply often exceeds the price. This provides an argument not for energy efficiency in general, but rather for load-shifting to off-peak periods, energy conservation during peak periods, or time-of-day pricing.

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 may 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 very uncertain and remote in time. With respect to retrofit investments, however, even this argument is unavailable, because these investments can simply be postponed until such time as they are needed to forestall the need to expand capacity.

4.4.2 Security Externalities

It has also been suggested that there are externalities associated with the economic and military security costs resulting from domestic U.S. dependence on oil imported 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.

5. NON-MARKET-FAILURE EXPLANATIONS OF THE ENERGY EFFICIENCY GAP

Non-market-failure explanations of the energy efficiency gap consist essentially of explaining why observed behavior is indeed optimal from the point of view of individual energy users. 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 calculations of what is optimal.

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 can quickly be concluded that no paradox exists, unless and until it can be demonstrated that the theoretical market failures already discussed are important in explaining observed behavior. This amounts to a policy prescription that markets should be "considered innocent until proved guilty." There is no scientific basis for such a prescription. On the other hand, it is important to keep in mind that although market failure may be considered a necessary condition for government policy intervention, it is not a sufficient condition. For government policy to be desirable in economic terms, 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. Analyses of past government and utility energy conservation programs demonstrate that although some such programs certainly have had positive net benefits, the benefits of others have been lower than predicted.

5.1 Irreversibility and the Option to Wait

One such non-market-failure explanation is that uncertainty about the future benefits of energy-efficient technologies, combined with the irreversible nature of the efficiency investment, makes the effective 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. Simple net-present-value analysis typically does not account for changes over time in the savings that purchasers might enjoy from an extra investment in energy efficiency, which depends on trends and uncertainties in the prices of energy and conservation technologies. When making irreversible investments that can be delayed, the presence of this uncertainty can lead to an investment hurdle rate that is larger than the discount rate used by an analyst who ignores this uncertainty.

The magnitude of this option-to-wait effect or option value depends on project-specific factors, such as the degree of energy price volatility, the degree of uncertainty in the cost of the investment, and how fast the prices of energy and conservation technologies are changing over time. The effect is magnified when energy and technology price uncertainty is increased and when energy prices are rising and technology costs are falling more quickly. Under conditions characterizing most energy conservation investments, this effect could raise the hurdle rate by up to about 10 percentage points. On the other hand, if there is no opportunity to wait, this effect can be ignored.

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 hurdle rates to irreversible investments when returns are uncertain. To the extent that consumers' true effective discount rates are high for this reason, this would not represent a market failure or a reason for policy intervention from an economic perspective.

5.2 Unaccounted for Costs and Overestimation of Energy Savings

There are typically costs of adoption that are not included in simple cost-effectiveness calculations. It is by no means costless to learn how a technological improvement fits into a home or firm or to learn about reliable suppliers. 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. Another type of hidden cost is 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 florescent and incandescent lighting and the delay time in achieving illumination with many florescent lights.

Some have argued that not only costly information acquisition but also biased estimates of likely energy savings play a role. There is evidence that analysts have substantially overestimated the energy savings that higher efficiency levels will bring, partly because projections often are based on highly controlled studies that do not necessarily apply to actual realized savings in a particular situation. For example, studies have found that actual savings from utility-sponsored programs typically achieve 50 to 80% of predicted savings. Other studies have drawn similar conclusions based on analysis of residential energy consumption data, finding that the actual internal rate of return to energy conservation investments in insulation was substantially below typical engineering estimates for the returns from such investments. On the other hand, the bias could be offset in the opposite direction, because some studies indicate that consumers systematically overestimate energy savings associated with some types of new technologies.

5.3 Heterogeneity in Energy Users

Another possible non-market-failure explanation for the energy efficiency gap is associated with the fact that even if a given technology is cost-effective on average, it will mostly likely not be for some individuals or firms. If the relevant population is heterogeneous—with respect to variables such as the purchaser's discount rate, the investment lifetime, the price of energy, the purchase price, and other costs—even a technology that looks very good for the average user will be unattractive for a portion of the population. Heterogeneity in these and other factors leads to differences in the expected value that individual purchasers will attach to more energy-efficient products. As a result, only purchasers for whom it is especially valuable may purchase a product. Depending on how the efficiency gap is measured, such underlying heterogeneity can provide yet another non-market-failure explanation. For example, it may not make sense for someone who will only rarely use an air conditioner to spend significantly more purchasing an energy-efficient model—they simply may not have adequate opportunity to recoup their investment through energy savings. Analysis based on single estimates for the important factors previously discussed—unless they are all very conservative—will inevitably lead to an optimal level of energy efficiency that is too high for some portion of purchasers. The size of this group, and the magnitude of the resulting inefficiency should they be constrained to choose products that are not right for them, will of course depend on the extent of heterogeneity in the population and the assumptions made by the analyst.

6. DISCOUNTING AND IMPLICIT 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 among a number of separate and distinct questions regarding the appropriate discount rate. First, an attempt can be made to estimate the implicit discount rate that consumers appear to be using when they make energy efficiency decisions. Second, there can be speculation about the correct discount rate for consumers to use in making such decisions. Third, the implicit discount rate (explained below) or the correct rate can be compared with other private discount rates in the economy. Fourth and finally, it can be asked whether any of these rates are equal to the social discount rate that ought to be applied in evaluating future energy savings when making public policy decisions.

The observation that consumers have high implicit discount rates when they make energy efficiency decisions is actually neither more nor less than a restatement of the existence of the energy efficiency gap. To estimate implicit discount rates, studies examine decisions actually made and calculate the discount rate that makes those decisions privately optimal, given 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. There is substantial empirical evidence of implicit discount rates for energy-conservation investment decisions in the range of 20 to over 100% for products such as air conditioners, space heating, and refrigerators. The energy efficiency gap 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 people make the decisions they make. 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 energy savings. Note that "actually" does 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. The truth is likely 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 only the relevant investment decisions are observed, it is fundamentally impossible to determine whether the observed behavior results from market failures, from truly high discount rates, or perhaps from misestimation of the true costs and savings of adoption. To make that distinction would require either observing 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 calculating from some basis other than the investment decisions what an appropriate discount rate for these investments would be. 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, postulating 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 it is found 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 that consumer worse off. 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. Instead, the question of whether relevant market failures exist is again the issue.

Finally, there is the issue of the social discount rate. Several decades of academic and policy debates have failed to resolve fully the question of what the social rate of discount should be, even as a conceptual matter. However, there is a strong argument that policies to increase energy efficiency can be justified only 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, and so on. It may be that the government should be doing more to encourage all of 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, then the set of available investment opportunities that should be undertaken probably exceeds the nation's annual gross national product. Obviously, the government can not undertake all of these projects. If the notion that public policy should have as its objective undertaking all investments for which rates of return exceed the social discount rate is taken seriously, then the reasonable prescription is to increase investment, beginning with the investments with the highest rates of return. But increasing investment reduces consumption, and long before investment opportunities were exhausted, consumption would be so significantly reduced that the social discount rate would rise. Increasing investment up to the point where the private and social rates of return coincide can be imagined, but it is impossible to know how much greater investment would be at that point or what the discount rate would be. 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, this form of investment could well be increasing at the expense of other forms 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.

7. SUMMARY

It is possible to synthesize much of the above discussion by examining graphically the concepts behind the major, alternative notions of the energy efficiency gap. To understand the basic elements of the debate, it is helpful to distinguish first between energy efficiency and economic efficiency, as in Fig. 3. The vertical axis measures increased energy efficiency (decreased energy use per unit of economic activity). The horizontal axis measures increased economic efficiency (decreased overall economic cost per unit of economic activity, taking into account energy and other opportunity costs of economic goods and services). Different points in the diagram represent the possible energy-using technologies available to the economy as indicated by their energy and economic efficiency. As a concrete illustration of this distinction, consider two air conditioners that are identical except that one has higher energy efficiency and, as a result, is more costly to manufacture because high-efficiency units require more cooling coils, a larger evaporator, and a larger condenser, as well as a research and development effort. Whether it makes sense for an individual consumer to invest in more energy efficiency depends on balancing the value of energy that will be saved against the increased purchase price, which depends on the value of the additional materials and labor that were spent to manufacture the high-efficiency unit. The value to society of saving energy should also include the value of reducing any associated environmental externalities; but again this must be weighed against the costs. Starting from the baseline situation, as represented by the lower left corner of Fig. 3, the adoption of more energy-efficient technology is represented as an upward movement. But not all such movements will also enhance economic efficiency. In some cases, it is possible simultaneously to increase energy efficiency and economic efficiency. This will be the case if there are market failures that impede the most efficient allocation of society's

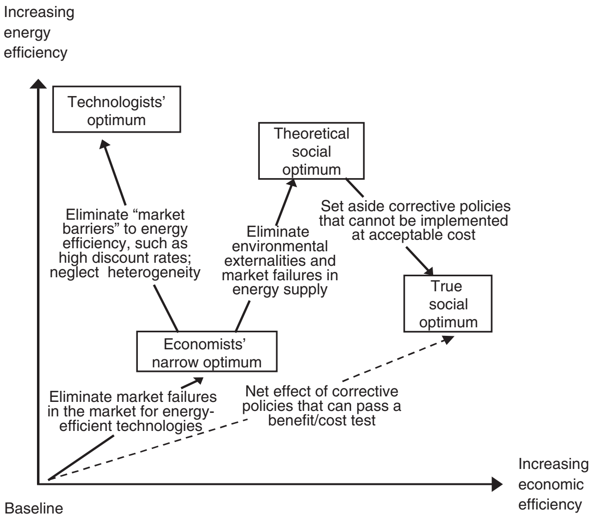

FIGURE 3 Alternative notions of the energy efficiency gap.

energy, capital, and knowledge resources in ways that also reduce energy efficiency. These are examples of what economists and others refer to as win-win or no regrets measures.

At the risk of excessive simplification, technologists can be characterized as believing that there are plentiful opportunities for low-cost or even negative-cost improvements in energy efficiency, and that realizing these opportunities will require active intervention in markets for energy-using equipment to help overcome barriers to the use of more efficient technologies. In terms of Fig. 3, the economist's notion of a narrow optimum is where market failures in the market for energy-efficient technologies have been corrected, the result being greater economic efficiency and energy efficiency. This optimum is narrow in the sense that it focuses solely on energy technology markets and does not consider possible failures in energy supply markets (such as underpriced energy due to subsidies or regulated markets) or, more important, environmental externalities associated with energy use (such as global climate change). Market failures in the choice of energy-efficient technologies could arise from a variety of sources, as discussed previously.

Eliminating broader market failures takes us to what is called the theoretical social optimum in Fig. 3. This represents both increased economic and increased energy efficiencies, compared with the economists' narrow optimum. This 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. But not all market failures can be eliminated at acceptable costs. In cases in which implementation costs outweigh the gains from corrective government intervention, it will be more efficient not to attempt to overcome particular market failures. This takes us from a theoretical social optimum to what is referred to as the true social optimum in Fig. 3. Market failures have been eliminated, but only those for which elimination can pass a reasonable benefit–cost test. The result is the highest possible level of economic efficiency, but a level of energy efficiency that is intermediate compared with what would be technologically possible.

In contrast to the economist's perspective, if non-market-failure market barriers (such as high discount rates caused by uncertainty about payback) were also removed, the technologists' optimum would be achieved. This alternative notion of an optimum is found by minimizing the total purchase and operating costs of an investment, with energy operating costs being discounted at a rate the analyst (not necessarily the purchaser) feels is appropriate. Clearly, if no distinction is made between barriers to adoption that are market failures and those that are not, and no choice is made to eliminate all barriers, a higher estimate of energy efficiency will be achieved, but not economic efficiency. An important implication of this perspective is that comparisons of an engineering ideal for a particular energy use with the average practice for existing technology are inherently misleading, because the engineering ideal does not incorporate all the real-world factors influencing energy technology decision-making.

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 is labeled technologists' optimum. 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 it is to specify what is the right approach. 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, the first requirement is to disentangle market-failure and non-market-failure explanations for observed decisions regarding energy-efficiency investments.

SEE ALSO THE FOLLOWING ARTICLES

Conservation Measures for Energy, History of • Consumption, Energy, and the Environment • Discount Rates and Energy Efficiency Gap • Economics of Energy Demand • Economics of Energy Supply • Energy Efficiency and Climate Change • Energy Efficiency, Taxonomic Overview • Industrial Energy Efficiency • Obstacles to Energy Efficiency

Further Reading

Anderson, S. T., and Newell, R. G. (2003). Information programs for technology adoption: The case of energy-efficiency audits. *Resource Energy Econ.* (in press).

DeCanio, S. J., and Watkins, W. E. (1998). Investment in energy efficiency: Do the characteristics of firms matter? *Rev. Econ. Statist.* **80**(1), 95–107.

Hassett, K. A., and Metcalf, G. E. (1993). Energy conservation investment: Do consumers discount the future correctly? *Energy Policy* **21**(6), 710–716.

Hassett, K. A., and Metcalf, G. E. (1995). Energy tax credits and residential conservation investment: Evidence from panel data. *J. Public Econ.* **57**, 201–217.

Hausman, J. A. (1979). Individual discount rates and the purchase and utilization of energy-using durables. *Bell J. Econ.* **10**, 33–54.

Hirst, E. (1986). Actual energy savings after retrofit: Electrically heated homes in the Pacific Northwest. *Energy* **11**, 299–308.

Jaffe, A. B., and Stavins, R. N. (1994). The energy paradox and the diffusion of conservation technology. *Resource Energy Econ.* **16**(2), 91–122.

Jaffe, A. B., and Stavins, R. N. (1994). The energy-efficiency gap: What does it mean? *Energy Policy* **22**(1), 804–810.

Jaffe, A. B., Newell, R. G., and Stavins, R. N. (2003). Technological change and the environment. *In* "Handbook of Environmental Economics," Vol. 1 (K.-G. Mäler and J. R. Vincent, Eds.), pp. 461–516. Elsevier, Amsterdam.

Joskow, P. L., and Marron, D. B. (1992). What does a negawatt really cost? Evidence from utility conservation programs. *Energy J.* **13**(4), 41–74.

Levine, M. D., Koomey, J. G., McMahon, J. E., and Sanstad, A. H. (1995). Energy efficiency policy and market failures. *In* "Annual Review of Energy and the Environment" (J. M. Hollander, Ed.), pp. 535–555. Annual Reviews Inc., Palo Alto.

Lind, R. C. (ed.). (1982). "Discounting for Time and Risk in Energy Policy." Resources for the Future, Washington, D.C.

Metcalf, G. E., and Hassett, K. A. (1997). Measuring the energy savings from home improvement investments: Evidence from monthly billing data. *Rev. Econ. Statist.* **81**(3), 516–528.

Newell, R. G., Jaffe, A. B., and Stavins, R. N. (1999). The induced innovation hypothesis and energy-saving technological change. *Q. J. Econ.*, August, 941–975.

Sanstad, A. H., Blumstein, C., and Stoft, S. E. (1995). How high are option values in energy-efficiency investments? *Energy Policy* **23**(9), 739–743.

Sebold, F. D., and Fox, E. W. (1985). Realized savings from residential conservation activity. *Energy J.* **6**(2), 73–88.

Sutherland, R. J. (1991). Market barriers to energy-efficiency investments. *Energy J.* **12**(3), 15–34.