

Technology Adoption and Aggregate Energy Efficiency

William A. Pizer, Winston Harrington,
Raymond J. Kopp, Richard D. Morgenstern and
Jhih-Shyang Shih

December 2002 • Discussion Paper 02–52



RESOURCES
FOR THE FUTURE

Resources for the Future
1616 P Street, NW
Washington, D.C. 20036
Telephone: 202–328–5000
Fax: 202–939–3460
Internet: <http://www.rff.org>

© 2002 Resources for the Future. All rights reserved. No portion of this paper may be reproduced without permission of the authors.

Discussion papers are research materials circulated by their authors for purposes of information and discussion. They have not necessarily undergone formal peer review or editorial treatment.

Technology Adoption and Aggregate Energy Efficiency

William A. Pizer, Winston Harrington, Raymond J. Kopp,
Richard D. Morgenstern, and Jhih-Shyang Shih

Abstract

Improved technology is often cited as a means to alter the otherwise difficult trade-off between the economic burden of regulation and environmental damage. Focusing on energy-saving technologies that mitigate the threat of climate change, we find that both energy prices and financial health influence technology adoption among a sample of industrial plants in four heavily polluting sectors. Based on a model linking technology adoption to growth in aggregate efficiency, we estimate that a doubling of energy prices, after raising the growth rate to 2.1%, would require slightly more than 50 years to generate a 50% improvement in aggregate efficiency relative to the baseline forecast.

Key Words: energy efficiency, endogenous technological change, technology adoption

JEL Classification Numbers: O31, O38, Q43, Q48

Table of Contents

1. Introduction.....	1
2. Background	2
3. Linking Technology Adoption and Aggregate Energy Efficiency.....	4
3.1. Aggregate Implications of Standard Diffusion Model.....	5
3.2. An Alternative Model of Technology Adoption and Energy Efficiency.....	7
3.3. Long-Run Behavior	9
3.4. Observable Variation	12
3.5. Aggregate Effects.....	13
4. Data and Estimation	15
4.1. Data on Technology Adoption and Plant Characteristics	15
4.2. Diffusion of Individual Technologies.....	17
4.3. Technology Adoption and Plant Characteristics.....	18
4.4. Prices, Profit, and Aggregate Energy Efficiency	20
5. Conclusion	22
References.....	32

List of Tables

Table 1: Descriptive Statistics in 1991	24
Table 2: Estimates of Diffusion Speed ^a	25
Table 3: Probit Model of Technology Adoption.....	26
Table 4: Changes in Plant and Firm Characteristics and Future Energy Efficiency	27

List of Figures

Figure 1: Effect of Diffusion Speed on Aggregate Energy Efficiency.....	28
Figure 2: Initial Distribution of Plant Efficiency and Efficiency Minus Adoption Tendency (simulated data).....	29
Figure 3: Distribution of Plant Efficiency and Efficiency Minus Adoption Tendency after 30 Years (simulated data)	30
Figure 4: Growth in Aggregate U.S. Energy Efficiency versus Energy Price Level.....	31

Technology Adoption and Aggregate Energy Efficiency

William A. Pizer, Winston Harrington, Raymond J. Kopp,
Richard D. Morgenstern, and Jhih-Shyang Shih*

1. Introduction

Many environmental problems pose a trade-off between environmental damage and expensive mitigation policies. This is particularly true of policies designed to mitigate global climate change, because carbon dioxide and other greenhouse gases remain in the atmosphere for centuries and mitigation is achieved primarily through reductions in the level and carbon intensity of energy use. Yet changes in technology could significantly alter these trade-offs. It is argued that public policies affecting the development and spread of new technologies may, over the long term, be one of the most important tools for environmental protection (Kneese and Schultz 1978).

But how can public policy encourage “environmentally-friendly” technological changes and how long will it take for them to substantially alter the mitigation cost–environmental damage trade-off? In the past, technology mandates have generally been used to deal with narrow environmental concerns within particular industries. However, the breadth of the problem faced by efforts to mitigate climate change and reduce the level and carbon intensity of energy use—energy used for a multitude of purposes across industries and time—defies such an approach. No one knows which technologies should be used in which industries, especially looking decades into the future. Instead, public policy must create incentives for economic agents to discover and adopt energy- and carbon-saving technologies of their own choosing.

To guide the design of incentive-based policies, this research examines some of the factors that influence the adoption of new energy-saving technologies by U.S. manufacturing

*The authors are Fellow, Senior Fellow, Senior Fellow, Senior Fellow, and Fellow, respectively, Quality of the Environment Division, Resources for the Future. The authors gratefully acknowledge financial support from the U.S. Department of Energy (Office of Energy Research, Integrated Assessment of Global Climate Change Research Program), and the U.S. Environmental Protection Agency (Office of Policy, Office of Economy and the Environment). Phillip Haberkern provided excellent research assistance. The research in this paper was conducted at the Center for Economic Studies, U.S. Bureau of the Census. Mary L. Streitwieser, Gordon M. Phillips, and Arnold P. Reznick have provided considerable assistance for which the authors are greatly appreciative. Research results are those of the authors and do not necessarily indicate concurrence by the Bureau of the Census or the Center for Economic Studies.

plants. We then take this analysis one step further and link changes in the rate of adoption to changes in the growth rate of aggregate energy efficiency. Using a unique dataset linking plant-level data from the Census Bureau's Longitudinal Research Database, technology use data from the Manufacturing Energy Consumption Survey, and parent-firm financial data from the Quarterly Financial Reports, we explore the adoption of four technologies in four industries. We analyze their adoption individually, focusing on diffusion speed, and then jointly, examining the influence of plant and firm characteristics on the adoption rate. Based on a model linking technology adoption to growth in aggregate energy efficiency, we extrapolate these results to determine the influence of these characteristics—primarily prices—on long-term energy efficiency improvements.

Our estimates of diffusion speed establish a remarkable consistency across technologies and industries. Once a technology has diffused to 10% of the plants, we estimate, the remainder of the plants will adopt it within an average of about nine years, regardless of the industry or the type of energy-saving technology.

When we consider the influence of plant and firm characteristics on the decision to adopt any new, energy-saving technology, we find that energy prices, plant size, and financial health have statistically significant effects. However, even dramatic changes in these variables generate only modest changes in aggregate energy efficiency for many years, based on the less invasive, incremental technologies that we examine. A doubling of energy prices, for example, requires 50 years to generate a 50% increase in energy efficiency over forecast levels at constant prices. The results also suggest that a policy to increase technology adoption through higher energy prices could backfire if financial health is compromised—a 50% reduction in profit, for example, more than offsets a 10% increase in energy prices in terms of the effect on technology adoption.

In the remainder of the paper, we start with a brief overview of the existing literature and theory on technology diffusion. In Section 3 we develop a simple model of technology adoption that establishes a link between plant-level adoption of energy-saving technologies and aggregate growth in energy efficiency. Section 4 presents our results and Section 5 concludes.

2. Background

Research by economists on the subject of technology diffusion dates back more than four decades. The single most important conclusion of previous work—well summarized by Jaffe and

Stavins (1994)—is that diffusion of new, economically superior technologies is a gradual rather than instantaneous process.¹ Specifically, diffusion is often portrayed as a classic S-shaped, or sigmoid, curve over time. That is, the rate of adoption begins slowly, speeds up, and then eventually slows down again as market saturation approaches.

One justification for the sigmoid curve is based on an epidemic model of diffusion. Because of lack of knowledge or confidence on the part of potential users, the odds that a nonuser will adopt a new technology increase with the growing popularity of the technology. If we let $N_{i,t}$ represent the presence of the new technology at time t for user i , we can write this relation as

$$P(\Delta N_{i,t} = 1 | N_{i,t-1} = 0) = 4c \cdot \bar{N}_{t-1} \quad (1)$$

where $\bar{N}_{i,t-1}$ is the average level of adoption—or popularity—at the beginning of period t and c is a constant (the “4” normalization is explained below). $P(\Delta N_{i,t} = 1 | N_{i,t-1} = 0)$ is the probability of a change in adoption status conditional on being a nonuser in the previous period. Since the likelihood of being a nonuser prior to time t is $(1 - \bar{N}_{i,t-1})$, the overall likelihood of a change in adoption status is

$$P(\Delta N_{i,t} = 1) = 4c \cdot \bar{N}_{t-1} \cdot (1 - \bar{N}_{t-1}) \quad (2)$$

Note that the parameter c reflects the maximum probability or speed of adoption that is obtained when the fraction of users $\bar{N}_{i,t-1}$ and nonusers $(1 - \bar{N}_{i,t-1})$ is each equal to 50%.

Following this intuition, it makes sense that the rate of adoption will be slow in the beginning ($\bar{N}_{i,t-1}$ is small when there is little popularity) and in the end ($1 - \bar{N}_{i,t-1}$ is small when there are few nonusers). In this model, the probability of adopting the technology depends entirely on the number of other firms in the industry that already have adopted it.

The pioneering work of Griliches (1957) extends this model by establishing that the diffusion of new technology can be understood in an economic framework by allowing the rate of diffusion to be partly determined by the expected economic return to adoption. Mansfield (1968) elaborates on this idea by considering the size of adopting plants, the perceived risk of the new technology, and the size of the required investment as potential determinants. The parameter

¹ See Griliches (1957), Mansfield (1968), David (1966), Davies (1979), Oster (1982), and Levin et al. (1987).

c in Equation (2) becomes a function of the economic return to adoption, primarily associated with technology characteristics but also determined by the broader environment for adoption.

The idea of a technology-specific diffusion speed has been explored by many authors. The size of potential users has been argued to have both positive and negative effects on adoption.² Arguments for the former are based on the resources (financial, experience, expertise) associated with large plants, and arguments for the latter hinge on potentially oligopolistic market structure retarding the competitive pressures to innovate. The possibility of varying diffusion rates for different technologies has been qualitatively described by Cohen and Levin (1989) as the difference between type “A” and “B” innovations.³ Type A innovations are minor and presumably diffuse quickly, whereas type B innovations are considerably more invasive and diffuse more slowly.

Looking beyond the aggregate determinates of adoption, David (1966) focuses on the idiosyncratic features of individual adopters via the use of more detailed microeconomic data. These features, which influence the value of the technology to individual plants, might include the cost of equipment, cost of learning about a new technology, cost of adapting existing processes, or future benefits of the technology. Specifically, one can imagine a threshold above which it pays to adopt the new technology and below which it does not. The threshold differs across plants and, over time, the cost of the innovation may fall and/or the quality may improve, thereby lowering the threshold. By regressing a plant’s decision to adopt on variables that describe differences among plants, one can empirically identify those differences that affect a plant’s valuation of the innovation.

3. Linking Technology Adoption and Aggregate Energy Efficiency

Like many of the aforementioned microeconomic studies, we use plant-level data on four incremental, energy-saving technologies to understand how both plant and technology characteristics influence the probability of adoption and speed of diffusion in four industries. However, we want to do more: We want to link these plant-level adoption decisions to macroeconomic changes in aggregate energy efficiency growth and, in the long run, to changes in aggregate energy use. Higher energy prices, for example, ought to encourage more rapid adoption of energy-saving technologies and lead to faster aggregate efficiency growth. Starting

² See Davies (1979), Oster (1982), and Boyd and Karlson (1993).

³ See also Davies (1979).

with microeconomic estimates that relate energy prices to technology adoption, if we can connect technology adoption to efficiency growth, we can then quantify the effect of energy prices on aggregate efficiency growth. To develop such a connection, we first consider the aggregate implications of a standard diffusion model, and then modify the model to permit this linkage.

3.1. Aggregate Implications of Standard Diffusion Model

When we consider the economy-wide impact of plant-level decisions to adopt energy-saving technologies, we can presume that any increase in the number of plant-level adoptions will, on average, raise aggregate energy efficiency. However, the diffusion model by itself fails to explain changes in the *long-term growth* in energy efficiency. Suppose, for example, that technology *X* has just begun to diffuse through the economy. Technology *Y* becomes available in 5 years and technology *Z* after 10. All three can raise aggregate efficiency by 10% once they are adopted by all plants. If each diffuses completely over 5 years, we would observe a 20% improvement after 10 years and 30% after 15 years. The solid line in Figure 1 shows graphically how adoption of these technologies would affect aggregate energy efficiency, assuming a constant diffusion rate.

Now imagine a policy that hastens the diffusion and leads to complete adoption in only 4 years for each technology, indicated by the dashed line in Figure 1. Note that this fails to change the long-term improvement in energy efficiency. Assuming technology *Y* still becomes available after 5 years and *Z* after 10 years, the total improvement in energy efficiency after 10 years is still 20%, and after 15 years 30%, even though the technologies diffuse more quickly in the first 4 years of each 5-year interval. The problem is that diffusion models focus on specific technologies but ignore the process by which those technologies become available—the process that eventually controls long-term growth in energy efficiency.

But consider what happens if we instead assume that technology options are always available: Just as one technology is adopted, another becomes available. In the previous example, imagine that technology *Y* becomes available after a 10% aggregate efficiency improvement and technology *Z* after a 20% improvement, regardless of when these improvements occur. Further, imagine that this pattern does not stop with technology *Z*; there is a never-ending stream of technologies that become available just as the previous ones are adopted. Perhaps technology purveyors foresee quicker returns when technologies diffuse faster, encouraging them to

innovate faster. Or perhaps the marketability of successive technologies simply requires adequate diffusion of the preceding technology.⁴ Although this view seems inappropriate for the large, radical, type B innovations discussed by Cohen and Levin (1989), it might be a reasonable approximation for smaller, frequent, and incremental type A technological improvements.

If we are willing to make this assumption, we can translate influences on plant-level adoption into persistent changes in the aggregate rate of technological change. For example, suppose that every year 100 plants adopt generic, energy-saving technologies, and that these adoption decisions are entirely responsible for an observed 1.3% annual growth rate in aggregate energy efficiency.⁵ Now suppose that a 10% rise in energy prices increases the number of annual adoptions from 100 to 107, a 7% rise. Based on the assumption that successive energy-saving technologies become available as soon as the preceding ones are adopted, this rise in the annual number of adoptions can persist as long as the incentive—higher energy prices—remains. We can then translate this 7% increase in adoptions into a 7% rise in the annual growth rate in aggregate energy efficiency—that is, from 1.3% to 1.4% ($= 1.07 \times 1.3\%$).

In this view of the world, the relationship presented in Equation (2) becomes inappropriate. To the extent we are concerned with improving aggregate efficiency, there is no need to evaluate or even identify individual technologies—we only need to track the generic decision to adopt new technologies each period. When a firm decides that its individual circumstances warrant an improvement in energy efficiency, it explores the menu of off-the-shelf technology options and chooses which ones to purchase.

Focusing on the decision of whether to invest each period in new energy-saving technologies, we can easily quantify the effect of different variables on the decision by using our cross-sectional dataset on plant-level technology choices. Our story linking plant-level adoption and aggregate energy efficiency then translates any estimated adoption effects into persistent effects on the annual growth rate of energy efficiency: More adoption leads to faster efficiency growth, and faster efficiency growth generates faster technology development, sustaining the initial increase in adoption and efficiency growth.

That link between technology adoption and aggregate growth in energy efficiency is exactly what we need so that we can use our microeconomic data on technology adoption to estimate effects on aggregate energy efficiency. Qualitatively, we can already see the necessary

⁴ Many authors have discussed the process of innovation in greater detail; see Jaffe et al. (2000).

⁵ Growth in aggregate energy efficiency averaged 1.3% between 1949 and 1999. See Section 4.4 and Figure 4.

assumptions. First, the efficiency improvement associated with each adoption decision needs to be constant unless we can predict how the size of discrete projects changes over time or in response to different incentives. Second, the availability of increasingly advanced technologies must be linked to the current technology level and divorced from calendar time—that way, new technologies are available as soon as the previous one is adopted. Third, we have to believe that technology adoption—specifically, adoption of incremental technologies of the sort contained in our dataset—is primarily responsible for aggregate efficiency growth.

Formalizing this story, we can develop insight about the precise assumptions that make the heuristic story work and about the dynamic link between plant-level adoption decisions and efficiency that we estimate. In particular, what do we need to assume about plant behavior and technology development to make sustainable increases in the adoption rate? How do we estimate and interpret an adoption model from a single cross section of plants?

3.2. An Alternative Model of Technology Adoption and Energy Efficiency

Suppose that each member i of a fixed population of I plants faces the decision of whether to invest in a discrete energy-saving project each period. The project permanently raises an adopting plant's energy efficiency (output per unit energy) by a fraction δ and incurs an annualized cost $J_{i,t}$. Assuming that this project is irreversible, $J_{i,t}$ should also incorporate the option value of waiting to invest. If $x_{i,t}$ is an index of plant i 's current efficiency, the plant should adopt if

$$\underbrace{Q_i \exp(-x_{i,t}) V_i}_{\text{current energy costs}} - \underbrace{Q_i \exp(-x_{i,t}) V_i \exp(-\delta)}_{\text{energy costs with energy-saving project}} > \underbrace{J_{i,t}}_{\text{cost of project}} \quad (3)$$

with

$$x_{i,t+1} = \begin{cases} x_{i,t} & \text{if plant } i \text{ does not adopt} \\ x_{i,t} + \delta & \text{if plant } i \text{ does adopt} \end{cases} \quad (4)$$

That is, the plant should adopt if the reduction in annual energy costs is greater than the annualized project cost, where Q_i is the output level of plant i , $\exp(x_{i,t})$ is the energy efficiency of plant i (output per unit of energy), and V_i is the energy price faced by plant i . So far, our model follows the more micro-oriented analyses of David (1966) and others, with adoption depending on differences among adopters.

The only controversial assumption is that the project size has been fixed for all plants and all time at δ . This is important: If we want to relate the plant-level adoption rate to growth in

aggregate energy efficiency, the size of the adopted technologies must remain constant—otherwise we have to consider changes to the size as well as the rate of adoptions. It seems reasonable to fix the size of the project as we consider the small, incremental, type A projects noted above. This is also a necessary assumption because we have little information about the size of the technologies in our dataset and therefore little ability to relax the assumption in practice.

Taking the log of both sides of (3) we can develop further assumptions about $J_{i,t}$. That is, firms adopt when

$$q_i + v_i + \log(1 - \exp(-\delta)) - x_{i,t} > j_{i,t} \quad (5)$$

where lowercase letters denote the logarithm of the corresponding uppercase variable— q_i is the log of the output level, v_i is the log of the energy price, $x_{i,t}$ is the log of the energy efficiency (which we will frequently refer to as simply “energy efficiency”), and $j_{i,t}$ is the log of adoption costs. From this relation, we can make the following observation: Our assumptions about $j_{i,t}$ will determine the growth rate of aggregate energy efficiency. If we assume that $j_{i,t}$ declines linearly over time, for example, then so will $x_{i,t}$. That is, $x_{i,t}$ needs to keep pace with $j_{i,t}$ for Equation (5) to remain roughly in balance and for optimal plant behavior to include both adoption and nonadoption over time.

To break the link between adoption costs and calendar time, thereby making the growth rate of energy efficiency endogenous, we assume that adoption costs $j_{i,t}$ decline one-for-one with \bar{x}_t , the average energy efficiency across all plants.⁶ This is the second controversial assumption we require, and it should not be surprising. We want to link a sustained increase in the adoption rate with a sustained increase in the growth rate of energy efficiency—regardless of the growth rate. Therefore, adoption costs need to be constantly matched to the current efficiency level because the efficiency level reflects the current energy cost savings from adoption. Adoption and efficiency growth cannot continue if adoption costs do not keep pace with the associated energy savings.

Consider how a $\delta = 10\%$ efficiency improvement saves less energy and less money as the absolute level of energy use and expense declines. The savings from technology adoption would be \$10 on a \$100 energy expenditure but only \$5 on a \$50 expenditure. When the average energy

⁶ An entirely different approach would be to assume that $j_{i,t}$ remains constant and that energy prices must rise to encourage further technology adoption.

bill is eventually reduced to \$50, thanks to improvements in energy efficiency, adoption costs need to roughly equal \$5 if adoption and efficiency improvements are going to continue—regardless of when the average energy expenditure reaches \$50. In this way, adoption costs must be linked to energy expenditures and/or the level of energy efficiency.

Based on the assumption that adoption costs fall one-for-one with mean efficiency \bar{x}_t , we write

$$j_{i,t} = j_i + \beta \cdot v_i + \gamma \cdot q_i + \varepsilon_{i,t} - \bar{x}_t \quad (6)$$

This further specifies that the remaining portion of adoption cost depends on plant output, $\gamma \cdot q_i$, the price of energy, $\beta \cdot v_i$, an idiosyncratic plant component, j_i , and a stochastic disturbance $\varepsilon_{i,t}$.

Common sense dictates that $0 < \beta < 1$.⁷ Note that with the introduction of a stochastic term in the cost of adoption and the dependence on average efficiency, the achievement of a particular level of efficiency x (e.g., technology diffusion) will occur gradually even when plants are identical. This incorporates the thinking in earlier epidemic models of technology diffusion that did not rely on heterogeneous adopters (Griliches 1957; and others).

From (6), the adoption condition becomes

$$(1-\gamma)q_i + (1-\beta)v_i + \log(1-\exp(-\delta)) - (x_{i,t} - \bar{x}_t) - j_i > \varepsilon_{i,t} \quad (7)$$

Assuming the stochastic element has a mean-zero normal distribution with variance s^2 , we can write the probability of adoption as

$$P(\text{adopt}) = \Phi \left(\frac{(1-\gamma)q_i + (1-\beta)v_i + \log(1-\exp(-\delta)) - (x_{i,t} - \bar{x}_t) - j_i}{s} \right) \quad (8)$$

where Φ is the cumulative standard normal density function.

3.3. Long-Run Behavior

With plant-level adoption decisions given by Equation (8), what sort of adoption behavior do we expect over time? Rewriting the adoption probability as

⁷ That is, higher energy prices encourage technology adoption; see Equation (7).

$$P(\text{adopt}) = \Phi\left(-\frac{x_{i,t} - g_i - \bar{x}_t}{s}\right) = \Phi\left(-\frac{(x_{i,t} - \bar{x}_t) - (g_i - \bar{g}) + \frac{\bar{g}}{s}}{s}\right) \quad (9)$$

and

$$g_i = (1 - \gamma)q_i + (1 - \beta)v_i + \log(1 - \exp(-\delta)) - j_i \quad (10)$$

where g_i summarizes a plant's exogenous adoption tendency, the last expression in Equation (9) suggests a convenient distinction between the mean \bar{g} across all plants and the deviations at a particular plant $g_i - \bar{g}$. The mean adoption tendency \bar{g} determines an overall adoption rate (the adoption rate for the mean plant), and the plant-level deviations both $g_i - \bar{g}$ and $x_{i,t} - \bar{x}_t$ determine the relative adoption likelihood for a specific plant.

Over time, we expect the difference between energy efficiency and adoption tendency $x_{i,t} - g_i$ to evolve to an ergodic distribution centered about \bar{x}_t . Plants with high adoption tendency ($g_i \gg \bar{g}$) and/or low initial efficiency ($x_{i,t} \ll \bar{x}_t$) will tend to adopt more rapidly, raising their value of $x_{i,t}$ relative to the mean, but plants with low adoption tendency ($g_i \ll \bar{g}$) and/or high initial efficiency ($x_{i,t} \gg \bar{x}_t$) will adopt less rapidly, lowering their value of $x_{i,t}$ relative to the mean. Absent any changes over time to the plant-level adoption tendency, g_i , this implies that eventually the adoption probability each period will be independent of these plant characteristics.

For example, imagine that there are two groups of plants, half with high adoption tendency, $g_i = g_{\text{hi}} = 0.2$, and half with low adoption tendency, $g_i = g_{\text{low}} = -0.2 < g_{\text{hi}}$ (so $\bar{g} = 0$). Let us assume that the initial energy efficiency $x_{i,0}$ is randomly distributed across all plants and is *initially* uncorrelated with adoption tendency. Figure 2 shows both the initial distribution of efficiency $x_{i,t}$ alongside the initial difference between efficiency and adoption tendency $x_{i,t} - g_i$ for this hypothetical population of plants. The pattern in the left panel, showing efficiency, reflects our assumption about the random distribution of efficiency levels, which is based on a mean-zero normal distribution with a standard deviation of 0.025. The right panel, showing the difference between adoption tendency and efficiency, reflects the assumption about two groups of plants—one characterized by high adoption tendencies (the left mode) and the other characterized by low adoption tendencies (the right mode). Here, the average distance between the two modes ($g_{\text{hi}} - g_{\text{low}}$) equals 0.4. The spread about each of these modes arises from the variation in $x_{i,t}$, so the standard deviation is once again 0.025.

Now imagine we allow time to pass. We let each period represent three years, corresponding to the three-year period in our data (1991 to 1994). The mean rate of adoption when $g_i = \bar{g}$ and $x_{i,t} = \bar{x}_t$ equals 50% over the three-year period, also matching our data (see last

column of Table 3). Given the initial conditions, however, we know that some plants have a higher probability of adoption than others: In the right panel of Figure 2 we have also plotted the adoption probability based on Equation (9) with an assumption that the standard deviation s of the stochastic disturbance $\varepsilon_{n,t}$ equals 0.1. Plants with high adoption tendency (left mode) adopt with an initial probability of about 98%, and plants with low adoption tendency (right mode) adopt with an initial probability of about 2%. Assuming a $\delta = 0.06$ improvement with each adoption—reflecting a three-year growth rate of 3% in energy efficiency when half the plants adopt—plants in the two modes will intermingle within seven periods because the initial variation in $x_{i,t} - g_i$ is only 0.4.

Figure 3 shows the results after 10 periods (30 years) of plant adoption behavior, simulated using Equations (9) and (4) with the parameter values given in the text: $\bar{g} = 0$, $\delta = 0.06$, and $s = 0.1$. In the left panel we see that plants with the higher adoption tendency have distinguished themselves from those plants with the lower adoption tendency by adopting more frequently and moving ahead in efficiency. Plants in the right mode, identified by the higher adoption tendency $g_{hi} = 0.2$, have an efficiency index $x_{i,t}$ that is about 0.34 higher on average than those plants in the left mode with $g_{low} = -0.2$. In another 30 years the difference in average efficiency $x_{i,t}$ between the two groups is 0.39, and in another 30 years after that the difference is 0.40—exactly the difference between g_{hi} and g_{low} . That is, the plants with a higher adoption tendency move ahead in energy efficiency just to the point where their higher relative efficiency offsets their naturally higher tendency to adopt. At this point, the likelihood of adoption is unrelated to the adoption tendencies measured by g_i .

The right panel again makes this point: The difference between efficiency and natural adoption tendency, $x_{i,t} - g_i$, is eventually indistinguishable across plants with different adoption tendencies. Because the adoption rate depends on this difference (as indicated by the overlaid plot of adoption probability), plants with lower values tend to adopt more rapidly and further raise their efficiency $x_{i,t}$ relative to the average. This brings them back toward the mean of the $x_{i,t} - g_i$ distribution and simultaneously slows their adoption rate by lowering their gain to further adoption. Here, the original difference between the high- and low-tendency plants visible in the right panel of Figure 2 is gone 10 periods later in the right panel of Figure 3.

This convergence in adoption rates among plants is not unlike the convergence in income levels across countries suggested by standard growth models (Barro 1991). In that case, higher returns to investment inspire faster capital accumulation in poorer countries, whereas here, higher returns to energy-efficient technologies inspire faster technology adoption. In both cases, as laggards converge to the mean, the incentive to catch up is diminished.

Although the cross-sectional relation between adoption rate and natural adoption tendency may disappear, this does not contradict a persistent *aggregate* relation between the average adoption rate and the average adoption tendency—the link that motivates our effort. A change in the average adoption tendency—a rise in the average value \bar{g} —will raise the adoption rate for all plants, a change that does not vanish with future adoptions. Put another way, the adoption rate for the average plant (where $x_{i,t} - g_i = \bar{x}_t - \bar{g}$) has permanently risen, a fact unaffected by any pattern of future adoptions based on Equation (9). The average plant will still observe an adoption rate given by $\Phi(\bar{g} / s)$.

3.4. Observable Variation

From the preceding discussion, it is apparent that variation in adoption rates across plants will be unrelated to any fixed plant characteristics g_i —eventually. For example, we would expect plants facing higher energy prices to be more likely to adopt energy-saving technologies, all other things equal. But if price variation across plants remains fixed over time, those plants facing higher energy prices will eventually move ahead in efficiency, reducing their gain from further adoption. That is, in terms of Equation (9), a plant's higher value of g_i due to higher prices will eventually be offset by a relatively higher value of $x_{i,t}$ and, without conditioning on $x_{i,t}$, the adoption rate will appear to be independent of prices.

This is problematic. The purpose of this paper is both to identify plant-level determinants of adoption and to extrapolate changes in adoption to changes in the aggregate growth rate in energy efficiency. This model, in seeking to do the latter, suggests that the former, simpler goal of estimating the determinants of adoption is in jeopardy. In essence, a plant that has a high propensity to adopt energy-saving technologies (because it faces relatively high energy prices, for example) may have just installed the preceding technology, making adoption unlikely in the current period. Despite the persistent relation between average adoption tendency \bar{g} and the average adoption rate $\Phi(\bar{g} / s)$, we will be unable to estimate the particular determinants of g_i from cross-sectional variation.

There are two ways to resolve this problem. First, plant characteristics are not completely fixed, despite the assumptions implied by Equation (3), where only the adoption cost and efficiency are explicitly time dependent. When we consider plant and firm characteristics in 1991—energy prices, employment, profit, and working capital, for example—it is easy to believe there have been plant-level changes in these variables in the not-too-distant past. Therefore, we are not yet in the long-run situation where plant characteristics are unrelated to plant adoption. Put another way, if we are closer to some initial period, such as the one depicted in Figure 2, we

would expect to see variation in adoption decisions based on plant characteristics. If we are closer to some later period, such as the one depicted in Figure 3, we would not.

Second, the eventual lack of a relationship between plant-level characteristics and adoption is due to a lack of conditioning on the current level of efficiency $x_{i,t}$ —in other words, an omitted variable bias. The direction of this bias will be negative because efficiency is positively correlated with adoption tendency and negatively correlated with the adoption rate. We can see this in the right panel of Figure 3, where the positive correlation of $x_{i,t}$ and g_i conceals the underlying distribution of g_i when viewing the related distribution of $x_{i,t} - g_i$. If we conditioned on $x_{i,t}$ and only considered plants with $x_{i,t} = 0.3$, for example, we would see a mode at $x_{i,t} - g_i = 0.5$ corresponding to $g_{\text{low}} = -0.2$ and a mode at $x_{i,t} - g_i = 0.1$ corresponding to $g_{\text{hi}} = 0.2$. Based on the plot of adoption probabilities, those with $g_{\text{low}} = -0.2$ would have virtually no chance of adoption, but those with $g_{\text{hi}} = 0.2$ would adopt with virtual certainty.

Unfortunately, we do not have good direct measures of efficiency $x_{i,t}$ to correct the potential bias. When we estimate our model, we will attempt to control for current efficiency using previous technology adoptions. However, their effectiveness as proxies also implies that these previous adoptions will be correlated with adoption tendency measured by the other predictors of adoption. This correlation suggests that the usefulness of the proxy in revealing the effect of other adoption predictors could be thwarted by problems with collinearity and that we should carefully interpret the results: If we are close to the steady state depicted in Figure 3, controls for previous technology adoptions ought to be helpful, but if we are close to the initial conditions depicted in Figure 2, these technology controls may confound our identification.

3.5. Aggregate Effects

With some idea of how to identify the model with cross-sectional data, we now consider how to extrapolate from our model of energy-saving technology adoption to growth in aggregate energy efficiency. Our adoption model connects the rate of technology adoption to changes in the average efficiency among the population of potential adopters. From Equation (4), we have

$$\bar{x}_{t+1} = \bar{x}_t + \delta \int_{-\infty}^{\infty} f_t(y) \cdot \Phi\left(-\frac{y}{s}\right) dy \quad (11)$$

where $y = x_{i,t} - g_i - \bar{x}_t$, $x_{i,t}$ is energy efficiency at plant i and time t , g_i is adoption tendency at plant i , \bar{x}_t is the mean efficiency at time t , Φ is the standard normal cumulative density, $\Phi(-y/s)$ is the probability of adoption from Equation (9), s is a scale parameter,⁸ and $f_i(y)$ is the density function of $(x_{i,t} - g_i - \bar{x}_t)$.⁹ The expression $\int_{-\infty}^{\infty} f_i(y) \cdot \Phi(-y/s) dy$ reflects the average adoption probability at time t , which for a density f_i that is relatively concentrated compared with the scale parameter s and not particularly skewed approximately equals $\Phi(\bar{g}/s)$. We can therefore approximate the growth in energy efficiency among our population of plants as

$$\mu = \bar{x}_{t+1} - \bar{x}_t \approx \delta \cdot \Phi(\bar{g}/s) \quad (12)$$

where μ is the periodic growth in energy efficiency. Note that any increase in $\Phi(\bar{g}/s)$ translates into a proportional increase in μ .

The further extrapolation from efficiency growth among our sample of manufacturing plants to economy-wide efficiency rests on two assumptions. We need to assume that adoption decisions within our sample of manufacturing plants are representative of energy-saving technology adoption decisions throughout the economy. We also need to assume that aggregate growth in energy efficiency is governed primarily by these kinds of adoption decisions.¹⁰ The first assumption is necessary because we have adoption data only on manufacturing plants and want to draw general conclusions—although there seems to be little reason to suspect that decisions at manufacturing plants are biased in one direction or another vis-à-vis the decisions made by other economic agents.

The second assumption is similarly necessary because we have little information concerning other sources of efficiency improvements. However, we have reason to suspect biases in this case: There are trends other than adoption that contribute to increasing energy efficiency at the aggregate level, such as a shift from manufacturing to service activities and the development and diffusion of larger, more radical technologies. To the extent these trends are exogenous to the policy changes we consider, we can view calculations based on our

⁸ The standard error of the random component of technology cost $j_{i,t}$ in (6).

⁹ Depicted in the right panels of Figure 2 and Figure 3. As discussed in Section 3.3, this evolves to an ergodic distribution.

¹⁰ One could also consider general equilibrium effects. For example, a reduction in energy use would lead to a reduction in the demand for intermediate goods to produce energy and a further indirect reduction in energy use per unit of final demand. This suggests that our partial equilibrium modeling would understate the aggregate effect.

assumptions as generous. That is, if we assume that all growth in energy efficiency is attributable to incremental technology adoption decisions, and then compute a fractional increase in growth based on an increase in the adoption rate, this will overstate the effect if some growth is exogenously related to other trends and unaffected by the increase in adoptions. On the other hand, if there are other endogenous trends, such as the adoption of larger, more radical, type B technologies, our analysis may understate the estimated effect on aggregate growth.

One thing that we can happily ignore is differences in energy usage across plants. For example, if large plants always adopted more frequently than small plants (for the moment assuming $\gamma < 1$ in (8)), we would be concerned that a simple measure of the average adoption rate would be a misleading indicator of changes in aggregate efficiency. The large plants should be counted more. However, we know that eventually such cross-sectional differences in the adoption rate cannot persist. Eventually large plants will reach a position of relatively high efficiency compared with small plants, at which point smaller savings per unit of output will exactly counteract the scale of effect of more output at large plants and they will adopt at the same rate as small plants.

4. Data and Estimation

So far, we have discussed models of both diffusion for one specific technology and adoption from a menu of technologies. We have referred to the fact that we have data on a cross section of plants indicating their possession of four different technologies at two points in time—and therefore their adoption decision during the intervening years. We now discuss the data in more detail and present our estimation results.

4.1. Data on Technology Adoption and Plant Characteristics

The data come from several large plant-level datasets collected periodically by the Census Bureau and the Department of Energy: the Longitudinal Research Database (LRD) (including data from the Annual Survey of Manufactures and the Census of Manufactures), the Manufacturing Energy Consumption Surveys (MECS), and the Quarterly Financial Report (QFR). Prior work by Morgenstern et al. (forthcoming) linked the first two of these to create a comprehensive database of operating characteristics over time for a large sample of manufacturing plants in four energy- and pollution-intensive industries: petroleum refining, plastics, pulp and paper, and steel. Long and Ravenscraft (1993) constructed a bridge linking these plant-level data to parent-firm financial data in QFR.

The 1991 and 1994 MECS data are of particular interest to this research. These surveys ask detailed questions about installed energy-saving technologies, both industry-specific and general, in the plant. In this study we focus on the four general energy-saving technologies available in all industries: computerized climate controls (HVAC), computerized process controls (process), waste-heat recycling (waste heat), and adjustable speed motors (motors).¹¹ Changes in the presence of these technologies between 1991 and 1994 are used to construct our measures of adoption.

After considering a variety of plant-level characteristics as predictors of adoption, we settled on two variables: employment—a proxy for plant size and output—and energy prices. Our basic adoption model (3) supports these choices. Further, no other plant-level variables provide robust results. The employment variable is based on the LRD measure of total employment in 1991 and logged for our analysis. The energy price variable is the state-level energy price index in 1991 for the state where the plant is located.¹²

One of the novel aspects of this research is our ability to consider firm-level characteristics that affect the adoption decision via the cost of adoption $J_{i,t}$ in Equation (3). Through the use of parent-firm identifying codes in LRD, we can link plant-level data with financial information from the parent firms in QFR (Long and Ravenscraft 1993). We explored a range of financial characteristics that we believed might affect the firm's cost of capital and, in turn, the cost of technology adoption. In the end, we settled on two: working capital and profitability. Working capital is defined as the excess of current assets over current liabilities, subtracting short-term debt and including installments on long-term debt. We scale this variable by total assets in our regression. Profitability is defined as the income (or loss) from operations

¹¹ Computerized climate controls allow the plant to reduce its energy bill by more efficiently controlling heating, ventilation, and air conditioning systems. Computerized process controls, in contrast, lead to indirect energy savings by improving quality and reducing waste. In the plastics industry, for example, computers are used to quickly recognize and remedy aberrations in the cooling conditions surrounding the raw plastic, reducing the amount of plastic that must be scrapped.

Waste heat recovery reduces energy use by substituting otherwise wasted heat for raw energy. In oil refining, for example, distilled oil is piped past incoming crude oil to preheat it, reducing the heat required for distillation. In steel, heat exchangers on smoke stacks are similarly used to channel steam back into the smelting process. Adjustable speed motors offer direct energy savings by eliminating the excess power provided by fixed-speed motors as needs fluctuate during production. Without this technology, this excess power is dissipated through friction or diverted to empty applications.

¹² The index is taken from EIA (1998). We experimented with price indices constructed from plant-level energy use data (expenditures divided by quantity), but this proved to be a very noisy measure. We also considered models with both logged prices and price changes; both led to weaker results.

divided by net sales. For both variables we construct a five-quarter average from the beginning of the second quarter of 1990 through the end of the second quarter of 1991.

Table 1 provides summary statistics for these variables in our sample.

4.2. Diffusion of Individual Technologies

Earlier discussions about diffusion established Equation (2) as a descriptive relation for the change in adoption status over time for a single technology. This differential equation can be solved for the path of adoption over time to yield

$$p(t) = \frac{\exp\left(4c \cdot \left(t - t_{\frac{1}{2}}\right)\right)}{1 + \exp\left(4c \cdot \left(t - t_{\frac{1}{2}}\right)\right)} \quad (13)$$

where $p(t)$ is the fraction of plants using the specified technology at time t and $t_{\frac{1}{2}}$ is a constant of integration and represents the point in time when one-half of the potential users have adopted the technology. Using (13) and our data indicating adoption status in 1991 and 1994, we can compute the parameter c , measuring the speed of diffusion. In particular, given $p(0) = p_0$ and $p(1) = p_1$, we can compute¹³

$$4c = \log(1/p_0 - 1) - \log(1/p_1 - 1) \quad (14)$$

Table 2 reports the diffusion speed c for each of the 16 industry-technology combinations available in our data. Standard errors are based on the sampling error in the measurement of p_0 and p_1 .¹⁴ Remarkably, the diffusion speed is essentially the same for all industry-technology combinations. That is, with 16 independent estimates of the diffusion speed, a chi-squared test of equality fails to reject at any level of significance. Given the three-year interval in the model (1991–1994), the estimated value, $c = 0.22$, indicates that the middle half of the population adopts the technology in about seven years.¹⁵ Put another way, after the technology has been adopted by 10% of the population, the average adoption time for the remaining 90% of the

¹³ If p_0 and p_1 are sufficiently close, this expression can be approximated by the ratio of the rate of adoption among nonusers, $P(\Delta N_{i,t} = 1 | N_{i,t-1} = 0)$, to the fraction of potential users that have already adopted, $\bar{N}_{i,t-1}$, based on (1).

¹⁴ The standard error of $p = x/n$ is given by $\sqrt{p(1-p)/n}$.

¹⁵ Note that $\frac{50\%}{22\%/\text{period}} \times 3(\text{years/period}) \approx 7 \text{ years}$. To get the exact value, we evaluate Equation (13).

population is about nine years.¹⁶ These estimates are consistent with previous studies of technology diffusion as well as estimates of capital depreciation.¹⁷

Of course, our real interest is in how various explanatory variables influence the adoption rate and growth in aggregate energy efficiency. At this point we could directly estimate the effect of our four predictor variables—energy prices, employment, profit, and working capital—on diffusion speed for each of the 16 industry-technology combinations. Doing so, however, would not serve our goals. We could conclude, for example, that doubling prices might lower the average adoption time noted above from nine to seven years, but this would not allow us to draw conclusions about growth in aggregate efficiency.¹⁸ As highlighted in Figure 1, raising the diffusion speed of one technology without affecting the availability of future technologies leaves long-run energy efficiency unchanged.

4.3. Technology Adoption and Plant Characteristics

We instead return to our model developed in Section 3. That is, we lump all the technologies together and consider how our predictor variables influence the choice to adopt any new energy-saving technologies over the period 1991–1994. The fact that the diffusion speed is the same across all four technologies (and all four industries) provides some support for our assumption that these technologies are similar.

Among plants that possessed three or fewer technologies in 1991—that is, they had the potential to adopt—we define our dependent variable y_i to equal one if at least one new

¹⁶ This is a trickier calculation of $\frac{1}{0.9} \int_t^{\infty} (s-t) dp(s)$ where t is the time when exactly 10% have adopted and $p(s)$ is the cumulative density function for adoption from Equation (13).

¹⁷ Mansfield (1961) presents 12 estimates of the speed parameter c in Table 1 of his paper. Converting from annual to 3-year rates (appropriate for our data) and correcting for the factor 4 in Equation (2), his average estimate is 0.22 (excluding tin containers, which had an outlying estimate nearly 10 times higher than any other value). Fraumeni (1997) reports a depreciation rate of 0.11 for industrial equipment. This suggests that after 14 years, newly invested equipment is worth only 20% of its original value. If equipment were typically replaced or improved at this point, and if there were an even distribution of installation dates, half the population would be installing new equipment—and potentially adopting new technologies—every 7 years.

¹⁸ Estimating the effect of these variables on diffusion speed, we find only 6 significant coefficients among the 64 industry-technology coefficient combinations—3 employment coefficients, 2 energy price coefficients (1 negative), and 1 profit coefficient. The average across technologies and industries of the energy price coefficient is 0.01 (with a standard error of 0.01). Using this estimate, a doubling of energy prices, from \$5.56 to \$10.12 per Btu, would raise the average diffusion speed from 0.22 to 0.27 and lower the average adoption time after 10% penetration to 7 years.

technology is acquired in 1994 at plant i , and zero otherwise. We then estimate via maximum likelihood the model

$$L(y_i, z_i) = (\Phi(\lambda'z_i))^{y_i} (1 - \Phi(\lambda'z_i))^{1-y_i} \quad (15)$$

where z_i is a vector of the predictor variables (employment, energy prices, profit, and working capital, defined above) for plant i , λ is a vector of coefficients, Φ is the standard normal cumulative density function, $\Phi(\lambda'z_i)$ is the probability of adoption, and $L(y_i, z_i)$ defines the likelihood function. Table 3 reports the coefficients normalized to measure their effect on the probability of adoption at the mean \bar{z} , e.g., $\phi(\lambda'\bar{z})\lambda$, where ϕ is the standard normal probability density function.

From Equations (9) and (10) and our discussion in Section , we know that the current level of energy efficiency, $x_{i,t}$, is an important control variable. Because we do not have a good direct measure of energy efficiency, we use the presence of energy efficiency technologies in 1991 as proxies and include them in the z_i vector: Plants that have recently adopted energy-saving technologies in the past should, other things equal, have higher efficiency. Because it is unclear how effective these proxies might be at explaining the relative differences in efficiency, however, we present estimates both with and without these controls in Table 3.

The sample size for individual industries tends to be small with correspondingly large standard errors. Only 2 or 3 of 16 coefficients (4 coefficients \times 4 industries) are significant, depending on the presence of technology controls. For that reason, we focus our discussion on the pooled estimates that combine observations across all 4 industries.¹⁹ In the pooled model, we find significant coefficients on employment and profit with technology controls, and significant coefficients on profit and energy prices without technology controls. With the exception of the coefficient on employment, the estimated coefficients are similar in magnitude and sign regardless of the presence of technology controls.

As discussed in Section , the desirability of technology controls hinges on whether the relative efficiency of different plants has equilibrated in response to each plant's natural adoption tendency. If relative efficiency has equilibrated, these controls are necessary to reveal any underlying relation. If not, their inclusion may introduce an unnecessary collinearity. In Table 3, we see these opposing effects. Employment, which is relatively constant for long periods and

¹⁹ Note that a chi-squared test of equality among coefficients across the four industries fails to reject at the 10% level for all 4 coefficients.

presumably has allowed equilibration, is insignificant unless technology controls are included (the value of 0.13 in the top panel is reduced to 0.01 in the bottom panel). On the other hand, coefficients on prices and financial characteristics, which change more frequently and are less likely to have allowed equilibration, change only slightly but lose some significance in the presence of technology controls.

In addition to their significance, the signs on each coefficient also coincide well with our theory. More working capital and profit, which should indicate improved financial health and lower financing and adoption costs, leads to increased adoption as predicted by the negative sign on adoption cost j_n in Equation (8). Higher energy prices also increase adoption, as predicted by the positive coefficient $1 - \beta$ on prices v_n in Equation (8). Although theory argues for both positive and negative size effects—the coefficient $1 - \gamma$ on output q_n is unsigned—we find positive effects. We cannot individually identify the underlying parameters γ , β , and s in the analytic model because we observe only the ratios $(1 - \gamma)/s$ and $(1 - \beta)/s$.

4.4. Prices, Profit, and Aggregate Energy Efficiency

Now that we have a general understanding of the pattern of signs and significance in Table 3, we can consider their quantitative implications for adoption and aggregate energy efficiency. We focus on prices and to a lesser extent profit because they are both potentially influenced by policy choices. Calculations for all variables are summarized in Table 4.

Working with the estimates without technology controls (bottom panel of Table 3), we find that the coefficient on energy prices is 0.06. From Table 1 we know that the mean energy price is \$5.56 per Btu, so a 10% increase in energy prices would raise the adoption rate at the mean plant by 3.3% ($\$0.56 \times 0.06$). That is, the mean adoption rate—approximately the adoption rate at the mean plant—would rise from 49.1% (last column of Table 3) to 52.4%.

Based on Equation (12), growth in aggregate energy efficiency is proportional to the mean adoption rate. If the adoption rate rises from 49.1% to 52.4%, then growth in energy efficiency should rise from 1.28% to 1.36%. That is, based on a historical growth rate of 1.28% per year,²⁰ we would expect a 10% increase in energy prices to raise this growth rate to $(0.524/0.491) \times 1.28 = 1.36$ % per year, a 0.08% increase. Note that a doubling of energy prices

²⁰ Between 1949 and 1999, aggregate energy efficiency rose from \$48.57 to \$91.60 per MBtu (constant 1996 dollars; EIA 2000). This reflects an annualized growth rate of 1.28%.

(a 100% increase) would have roughly 10 times the effect, raising growth in energy efficiency to 2.05%. Column (6) of Table 4 summarizes these growth calculations.

The estimates of price effects match up well with estimates based on aggregate data, despite our extrapolation from individual plant behavior and heroic assumptions about innovation, technology costs, and the growth in aggregate efficiency. Figure 4 shows the growth in energy efficiency (averaged over 5 years) alongside energy prices since 1970. The doubling of energy prices from 1970 to 1980 was associated with a 2–3% growth rate in energy efficiency during the early 1980s versus an average of 1.3% over the past 50 years. A regression of aggregate efficiency growth on energy prices suggests that a doubling of energy prices would raise efficiency by 1.4% (with a standard error of 0.7%), from 1.3% to 2.7%.²¹ Both of these observations are consistent with our plant-level estimates.

In addition to statistical significance and aggregate consistency, we can also ask about the practical significance of these estimates. In particular, how much will these changes in the growth rate affect future energy efficiency? Consider that energy efficiency would rise from \$92 per MBtu today to \$173 per MBtu by 2050 based on an extrapolation of the historical rate of energy efficiency growth, 1.28%. Raising that growth rate to 1.36% via a 10% increase in energy prices would imply that energy efficiency reaches \$180 per MBtu by 2050—an improvement over the baseline energy efficiency of only 4.3% after 50 years. However, doubling energy prices and raising the growth rate to 2.05% would achieve an efficiency of \$253 per MBtu by 2050, a 46% improvement over the baseline. These calculations are summarized in columns (7) and (8) of Table 4.

All of this suggests that small increases in the price of energy will have marginal consequences for future efficiency, and even large price increases will require years to have a significant impact. Such results are consistent with our historical observations of gradual efficiency improvements—even during the energy price shocks of 1970s, we did not see dramatic, rapid reductions in energy use.

A final cautionary note arises from the consequences associated with profit. Applying the above calculations to the profit coefficient, we find that a 10% increase in profit, from a rate of 0.045 to 0.050, would raise the adoption rate from 49.1% to 49.9% based on the parameter

²¹ Based on a simple regression of efficiency growth on energy prices that allows for first-order autocorrelation of the errors, the coefficient on energy prices is 1.35 with a standard error of 0.65 (the estimated error autocorrelation of 0.73 is also significant).

estimate of 1.72 without technology controls. This would correspond to an increase in the growth rate of energy efficiency from 1.28% to 1.30%. The cautionary note arises if we believe that increased energy prices might lower profit. If, for example, a 10% increase in energy prices cannot be passed on to consumers, this might lower profit—perhaps by more than 10%.²² For example, a 50% decline in profits would more than offset the effect of a 10% rise in energy prices and lead to a decline in adoption and the growth rate of aggregate energy efficiency, rather than an increase.

5. Conclusion

From those results we draw two qualitative conclusions. Because we find that financial health has a significant effect on adoption, it will be important to carefully consider the broader economic impact of public policy to spur greater energy efficiency. A policy that immediately raises energy prices without allowing firms to anticipate the change may lead to a declining adoption rate if their financial health is adversely affected. Second and more significantly, the results suggest that improved technology will not generate significant improvements in energy efficiency for many years. Because improvements in energy efficiency currently progress at only 1.3% per year, even a large increase in this growth rate will take many years to translate into significant reductions in energy use below the baseline forecast.

Both of those conclusions follow naturally from what we know about past trends in aggregate energy efficiency as well as what we can observe in a simple model relating adoption to plant and firm characteristics. A novel part of this paper, however, is a model of technology adoption that quantitatively links individual plant behavior to aggregate energy efficiency. We can then estimate the actual change in energy efficiency growth arising from higher energy prices. Doing so, we find that for every 10% rise in prices, the annual growth rate of energy efficiency rises by 0.08%—that is, from 1.28% to 1.36%. A 10% increase itself leads to only marginal reductions in energy use even after 50 years. A doubling of energy prices, however, generates about a 2% growth rate in energy efficiency and an overall increase in efficiency of about 50% after 50 years (a one-third reduction in energy use).

²² The 1997 share of energy expenditures in total costs were, 0.05, 0.05, 0.02, and 0.10 in pulp and paper, plastics, petroleum, and steel, respectively (U.S. Department of Commerce 1997). Given that the mean profit rate was 0.05, this suggests that a 10% rise in energy prices would lower profit by 5% to 20%, ignoring any demand effects.

This research casts doubt on the idea that modest incentives can lead to large energy savings but demonstrates that more dramatic incentives—such as a doubling of energy prices—can eventually generate a substantial effect. Yet there are important caveats. We have focused on incremental technologies that do not fundamentally change the way plants operate—type A improvements in the literature. It is entirely possible that larger, type B innovations could be encouraged along with the incremental technologies, and that these larger innovations could offer substantial improvements in the nearer term. A second possibility is that aggregate changes in energy prices may bring forth faster technological development than revealed by our cross-sectional analysis, a problem inherent whenever cross-sectional data are applied to aggregate questions. Finally, some of our assumptions linking technology adoption to aggregate energy efficiency (e.g., the constant size of the efficiency gain) could prove to be too conservative. Despite our conclusions that large incentives and long horizons are necessary to substantially alter aggregate energy efficiency, these caveats suggest ways that the economy-versus-the-environment trade-off could be altered even in the near term, as well as indicate potentially fruitful areas for further research.

Table 1: Descriptive Statistics in 1991

	<i>Mean</i>	<i>Standard deviation</i>
Employment (logged)	6.050	1.085
Energy prices (\$ per MBtu)	5.563	1.265
Working capital (fraction of total assets)	0.085	0.116
Profitability (fraction of net sales)	0.045	0.055
<i>Dummy variables</i>		
Installed computer HVAC	0.165	
Installed computer process controls	0.661	
Installed waste heat recycling	0.579	
Installed variable-speed motors	0.589	
Pulp and paper plant	0.399	
Petroleum refinery	0.231	
Steel mill	0.174	
Plastics plant	0.196	
Total plants	316	

Table 2: Estimates of Diffusion Speed^a
(standard errors in parentheses)

	<i>Pulp and paper</i>	<i>Steel</i>	<i>Plastics</i>	<i>Petroleum</i>
HVAC	0.25 (0.05)	0.12 (0.06)	0.22 (0.06)	0.24 (0.06)
Process	0.26 (0.05)	0.42 (0.10)	0.29 (0.08)	0.23 (0.06)
Waste heat	0.15 (0.04)	0.17 (0.05)	0.20 (0.05)	0.24 (0.08)
Motors	0.30 (0.07)	0.33 (0.07)	0.17 (0.06)	0.24 (0.05)
Plants (n)	126	55	62	73

Precision weighted average:0.22 (0.01)

Simple average:0.24 (0.02)

Chi-squared test:18.2 (*p*-value of 0.31)

(all equal to 0.22; $\sim \chi^2$ (16))

^aDiffusion speed is the fraction of the population adopting a technology each period at the diffusion midpoint (50% adoption).

Table 3: Probit Model of Technology Adoption
 (dependent variable is adoption of one or more new technologies during 1991-1994)

	<i>increase in adoption probability for a unit change in:</i>					
	<i>working capital</i>	<i>energy prices</i>	<i>employment</i>	<i>profit</i>	<i>sample size</i>	<i>mean adoption rate</i>
	<i>(% of assets)</i>	<i>(\$/btu)</i>	<i>(logged)</i>	<i>(% of sales)</i>		
<i>with 1991 technology controls:</i>						
pulp and paper	1.45 (0.85)	0.05 (0.05)	0.05 (0.06)	2.03 (1.15)	112	0.455
plastics	0.16 (0.68)	-0.05 (0.07)	0.17 (0.12)	5.67* (2.21)	56	0.500
petroleum	-0.90 (0.68)	0.07 (0.06)	0.17* (0.08)	-0.97 (1.58)	66	0.500
steel	1.46* (0.70)	-0.02 (0.11)	0.11 (0.10)	0.42 (1.75)	51	0.549
pooled	0.23 (0.31)	0.04 (0.03)	0.13* (0.04)	1.54* (0.65)	285	0.491
<i>without 1991 technology controls:</i>						
pulp and paper	1.19 (0.74)	0.09* (0.04)	-0.05 (0.05)	1.99 (1.04)	112	0.455
plastics	-0.21 (0.57)	-0.02 (0.06)	0.03 (0.09)	3.69* (1.57)	56	0.500
petroleum	-0.45 (0.62)	0.07 (0.05)	0.07 (0.06)	-0.61 (1.51)	66	0.500
steel	1.05 (0.56)	-0.04 (0.09)	-0.01 (0.07)	0.98 (1.43)	51	0.549
pooled	0.26 (0.28)	0.06* (0.03)	0.01 (0.03)	1.72* (0.60)	285	0.491

*Asterisks indicate significance at the 5% level.

^aTechnology controls reflect the presence of four dummy variables indicating whether or not each of the four technologies was installed in 1991.

^bPooled model includes dummy variables for each industry (not reported).

Table 4: Changes in Plant and Firm Characteristics and Future Energy Efficiency

	<i>Change^b</i>	<i>Estimated coefficient^a</i>	<i>Change in adoption rate (%)</i>		<i>New ÷ old adoption rate^d</i>	<i>Growth in energy efficiency (%)</i>	<i>\$GDP per MBtu after 50 years</i>	<i>Relative to baseline^e (%)</i>
	<i>(1)</i>	<i>(2)</i>	<i>Approximate^c</i>	<i>Exact^c</i>	<i>(5)</i>	<i>(6)</i>	<i>(7)</i>	<i>(8)</i>
<i>Baseline</i>						1.28	173	
<i>10% increase</i>								
Working capital	0.01	0.26	0.2	0.2	1.00	1.29	174	0.3
Energy prices	0.56	0.06	3.3	3.3	1.07	1.36	180	4.3
Employment	0.61	0.13	7.7	7.7	1.16	1.48	191	10.4
Profit	0.005	1.72	0.8	0.8	1.02	1.30	175	1.0
<i>100% increase</i>								
Working capital	0.09	0.26	2.2	2.2	1.04	1.34	178	2.9
Energy prices	5.56	0.06	32.5	29.5	1.60	2.05	253	46.1
Employment	6.05	0.13	77.2	48.2	1.98	2.54	321	85.5
Profit	0.05	1.72	7.7	7.7	1.16	1.48	191	10.4

^aColumn 1 indicates the estimated coefficients from the pooled model without technology controls, bottom panel of Table 3, except employment.

^bColumn 2 indicates the change in the corresponding right variable, computed from the mean reported in Table 1.

^cApproximate indicates the product of the variable coefficient multiplied by the specified change, approximating the change in the adoption probability. Exact indicates the exact change in adoption probability for the mean plant, taking into account the curvature of the probit function.

^dColumn 5 is computed as $0.49 + (\text{previous column}) \div 0.49$, where 0.49 is the adoption rate of the mean plant.

^eThe baseline level of energy efficiency after 50 years is \$173 of GDP per MBtu (based on a growth rate of 1.28%).

Figure 1: Effect of Diffusion Speed on Aggregate Energy Efficiency

(dashed line reflects faster diffusion)




Figure 2: Initial Distribution of Plant Efficiency and Efficiency Minus Adoption Tendency (simulated data)

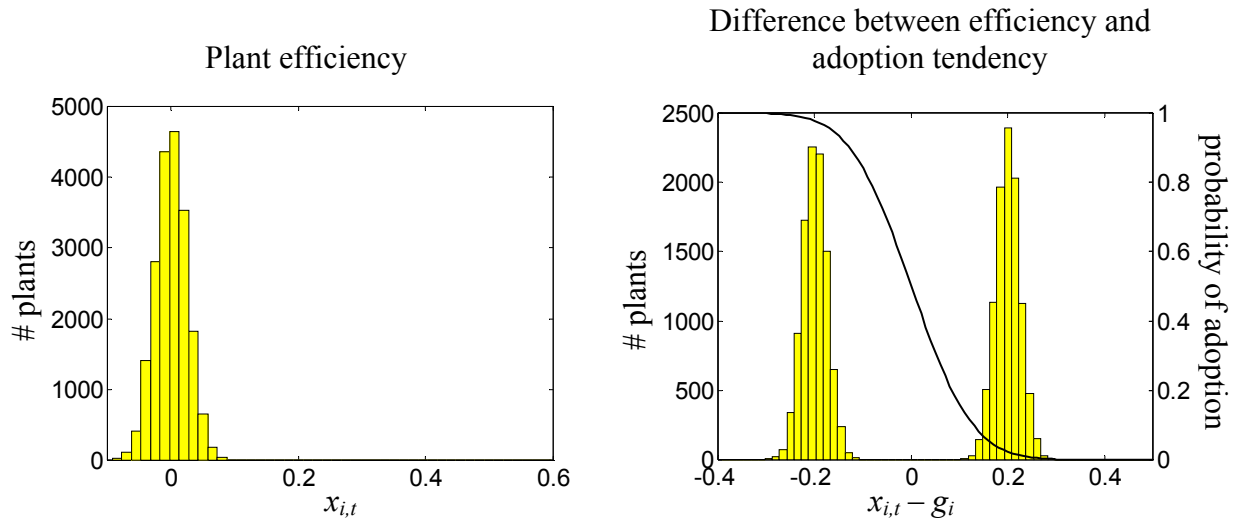


Figure 3: Distribution of Plant Efficiency and Efficiency Minus Adoption Tendency after 30 Years (simulated data)

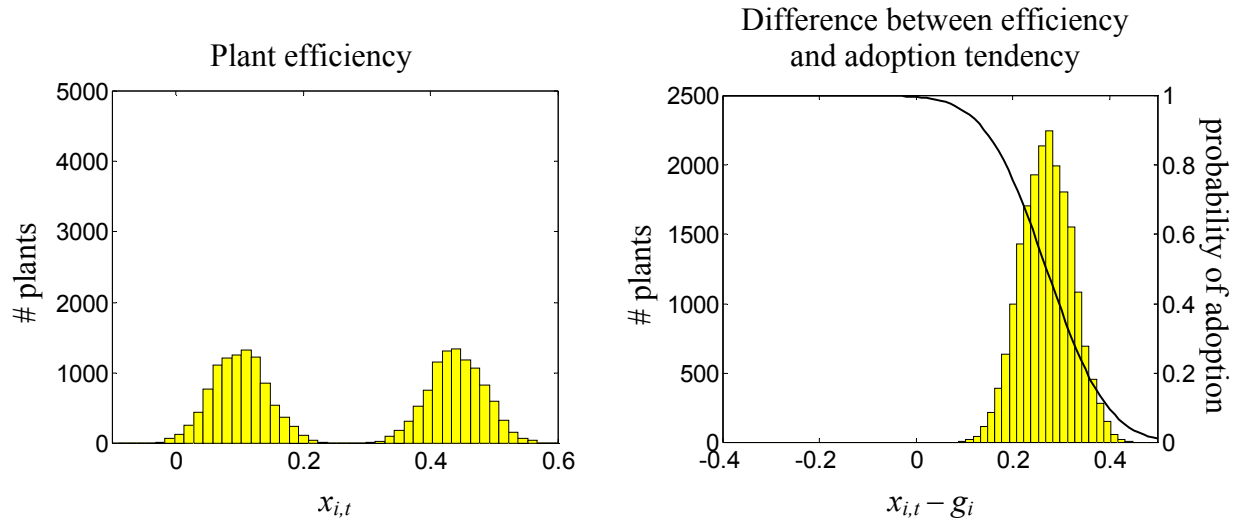



Figure 4: Growth in Aggregate U.S. Energy Efficiency versus Energy Price Level



Source: Table 1.5 (growth in energy efficiency) and Table 3.3 (energy prices), EIA (2000); Table B-7 (GDP price index), CEA (2000).

References

- Barro, Robert J. 1991. Economic Growth in a Cross Section of Countries. *Quarterly Journal of Economics* 106(2): 407–43.
- Boyd, Gale, and Stephen Karlson. 1993. The Impact of Energy Prices on Technology Choice in the US Steel Industry. *Energy Journal* 14(2): 47–56.
- Cohen, Wesley, and Richard Levin. 1989. Empirical Studies of Innovation and Market Structure. In *Handbook of Industrial Organization*.
- Council of Economic Advisors (CEA). 2000. *Economic Report of the President Transmitted to the Congress*. Washington: GPO.
- David, P.A. 1966. The Mechanization of Reaping in the Ante-Bellum Midwest. In *Industrialization in Two Systems*, edited by H. Rosovsky. Cambridge, MA: Harvard University Press.
- Davies, S.W. 1979. Inter-firm Diffusion of Process Innovation. *European Economic Review* 12: 299–317.
- Energy Information Administration (EIA). 1998. State Energy Price and Expenditure Report 1995. Washington, DC: U.S. Department of Energy/EIA. 0376(95).
- . 2000. Annual Energy Review 1999. Washington, DC: U.S. Department of Energy/EIA. 0384(99).
- Fraumeni, Barbara M. 1997. The Measurement of Depreciation in the U.S. National Income and Product Accounts. *Survey of Current Business* 77(7): 7–23.
- Griliches, Z. 1957. Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica* 25: 501–22.
- Jaffe, A.B., and R.N. Stavins. 1994. The Energy Paradox and the Diffusion of Conservation Technology. *Resource and Energy Economics* 16: 91–122.
- Jaffe, Adam B., Richard G. Newell, and Robert N. Stavins. 2000. Induced Invention, Innovation, and Diffusion: An Integrated Application to Energy-Saving Technology. Washington, DC: Resources for the Future. Working Paper.
- Kneese, Allen V., and Charles L. Schultz. 1978. *Pollution, Prices and Public Policy*. Washington, DC: Brookings Institute.

- Levin, S.G., S.L. Levin, and J.B. Meisel. 1987. A Dynamic Analysis of the Adoption of a New Technology: The Case of Optical Scanners. *Review of Economics and Statistics* 69: 12–17.
- Long, W.F., and D.J. Ravenscraft. 1993. The Quarterly Financial Report Database.
- Mansfield, Edwin. 1961. Technical Change and the Rate of Imitation. *Econometrica* 29(4): 741–65.
- . 1968. *Industrial Research and Technological Innovation*. New York: W.W. Norton.
- Morgenstern, Richard D., William A. Pizer, and Jhih-Shyang Shih. Forthcoming. The Cost of Environmental Protection. *Review of Economics and Statistics*.
- Oster, S. 1982. The Diffusion of Innovation among Steel Firms: The Basic Oxygen Furnace. *The Bell Journal of Economics* 13: 45–56.
- U.S. Department of Commerce, Bureau of Economic Analysis. 1997. Benchmark Input-Output Tables. *Survey of Current Business*.