MIT

CEEPR   Center for Energy and Environmental Policy Research

Energy Prices and the Adoption of
Energy-Saving Technology

by
Joshua Linn

06-012 WP   April 2006

A Joint Center of the Department of Economics, Laboratory for Energy
and the Environment, and Sloan School of Management

Energy Prices and the Adoption of Energy-Saving Technology *

Joshua Linn

April 2006

Abstract

This paper investigates the link between factor prices, technology and factor demands. I estimate the effect of price-induced technology adoption on energy demand in the U.S. manufacturing sector, using plant data from the Census of Manufactures, 1963-1997. I compare the energy efficiency of entrants and incumbents to measure the effect of technology adoption on the demand for energy. A 10 percent increase in the price of energy causes technology adoption that reduces the energy demand of entrants by 1 percent. This elasticity has two implications: first, technology adoption explains a statistically significant but relatively small fraction of changes in energy demand in the 1970s and 1980s; and second, technology adoption can reduce the long run effect of energy prices on growth, but by less than previous research has found.

Keywords: Constant Elasticity of Substitution, Endogenous Technological Change, Energy Prices, Long Run Growth, Manufacturing, Technology Adoption
JEL Classifications: Q41, Q43

1 INTRODUCTION

This paper investigates the link between factor prices, technology and factor demands. I estimate the effect of price-induced technology adoption on energy demand in the U.S. manufacturing sector, and find that an increase in the price of energy causes technology adoption that significantly reduces demand.

Technology adoption plays a central role in the relationship between energy prices and long run growth. Amid recent concerns that rising oil prices may harm the U.S. economy, former Federal Reserve Chairman Alan Greenspan predicted that, "Unless oil prices fall back, some of the more oil-intensive parts of our capital stock would ... be displaced, as was the case following the price increases of the late 1970s."[1] Many economists have argued that energy prices have a negative effect on GDP in the short run, but that technological change reduces energy demand and allows GDP to recover in the long run.

Atkeson and Kehoe (1999) make this argument formally. In their model, technology is "putty-clay"; a plant chooses the capital to energy ratio of its machines, and cannot adjust it later. Because the capital to energy ratio is fixed in the short run, energy demand is inelastic, and a price increase causes a large decrease in output. In the long run, plants choose a higher capital to energy ratio, reducing energy demand and increasing output.

In accordance with this model's predictions, total energy demand fell soon after the price increases in the 1970s, and output recovered. For two reasons, however, there is little direct evidence that technological change explains these patterns. First, the oil price increases were aggregate shocks, making it impossible to infer causality from a purely time series relationship between the price of energy and energy demand.

Second, aggregate data does not distinguish substitution along the energy demand curve from a shift of the curve due to technological change. Figure 1 shows this ambiguity, plotting the price of energy and energy demand, $D(A, \sigma)$, where A is the level of technology and σ is the energy price elasticity of demand for the particular technology. The economy begins at the point (E_0, P_0) and when the price increases to P_1, energy demand falls to E_1. If the initial energy demand curve were $D^\sigma(A_0, \sigma_1)$, plants would substitute away from energy and move along the curve. Alternatively, substitution possibilities may be limited with the technology, and the initial demand curve would be $D(A_0, \sigma_0)$. To reach the point (E_1, P_1), firms would adopt technology A_1, and the demand curve would shift to $D^A(A_1, \sigma_0)$. Thus, firms may have moved to the point (E_1, P_1) either by substituting or by adopting new technology. It is necessary to estimate the initial demand curve for energy to determine whether technology adoption has occurred.

---

[1]Remarks to the National Italian American Foundation, Washington, D.C., October 15, 2004.

The empirical strategy employs two techniques to address these issues, allowing me to estimate the effect of price-induced technology adoption on energy demand. First, the use of plant level data allows me to exploit cross sectional variation in the effect of an aggregate shock. Because industries vary in their energy intensity, a given shock affects industries differently.

Second, I compare the energy efficiency of entering and existing plants. Assuming that existing plants cannot adopt technology, they move along their demand curve when the price of energy increases. The comparison between incumbents and entrants reveals the effect of technology adoption because if entering plants use the same technology, they would have the same energy efficiency as incumbents. If they use less energy than existing plants, this difference would correspond to an inward shift of their demand curve.[2]

I find that a 10 percent increase in the price of energy leads to technology adoption that causes a one percent decrease in energy demand. The real price of energy rose by a factor of about 2.5 between 1972 and 1982, so that, relative to the difference in 1972, entrants in 1982 were 25 percent more efficient than incumbents.

The results imply that technology adoption is of lesser importance than theoretical and anecdotal evidence suggest. Substitution and changes in industry composition explain a larger fraction of the historical reduction in energy use in U.S. manufacturing.[3] Technology adoption reduces energy prices' long run effect on output by a relatively small amount.

I now discuss the identification strategy in more detail. Energy prices are potentially endogenous; for example, technological change could reduce demand and lower prices. To address this concern I construct a fixed-weight index of the total price of energy. The index uses the sector-wide average prices of different fuels and electricity, and should not include the effects of technological change, substitution or industry-specific shocks to output demand.[4]

The estimating equation rests upon three assumptions. First, the total factor productivity (TFP) differential between entering and existing plants is uncorrelated with the price of energy. Second, the elasticity of substitution is constant and equal for entrants and incumbents. Third, existing plants do not adopt new technology. While these assumptions are strong, they allow me to estimate a linear equation and compare the response of incumbents and entrants to

---

[2]As I discuss below, to the extent that incumbents adopt technology, my results underestimate the total effect of technology adoption. Note that both the adoption of newly invented machines and previously existing ones could cause the demand curve to shift.

[3]Industry composition has not been particularly important since about 1980, but explained much of the decline in total energy demand before 1980 (Wing and Eckaus, 2004).

[4]Another issue related to the price of energy is the fact that in most of the analysis I use the current price of energy as the independent variable. I assume that the current price captures the permanent component of a price shock; this would be the case if the price follows a random walk. Below I relax this assumption, and compute a forecasted price of energy, which yields similar results. Note that if the price of energy is mean-reverting, the results underestimate the actual effect of technology adoption.

energy prices. In the remainder of the empirical analysis, I document that the first and third assumptions seem to hold in practice, and that relaxing the second assumption does not affect the results.

More specifically, a plant's TFP affects its energy demand. If an increase in the price of energy causes the entry of plants with higher than average TFP (as in the model discussed below) the estimate would be biased away from zero. That is, when the price is high, entrants would be more efficient because they require less of all inputs, but otherwise use the same technology. Several empirical approaches suggest this is not a major concern.

Second, regarding the substitution elasticity of entrants and incumbents, Diamond, McFadden and Rodriguez (1978) argue that it is difficult to empirically distinguish technological change from a change in the substitution elasticity. For example, consider Figure 1, and assume that plants initially operate on the demand curve $D(A_0, \sigma_0)$. After a price shock entrants can either be on the demand curve $D^\sigma(A_0, \sigma_1)$, which has a greater price elasticity than the original technology, or on the curve $D^A(A_1, \sigma_0)$, which is an inward shift of the original curve. It appears that the baseline estimate reflects the latter case; the short run price elasticity, σ, is similar for entrants and incumbents.

Finally, I assume that existing plants do not adopt technology. To the extent that this assumption does not hold, the results would underestimate the total effect of technology adoption. As documented below, the average incumbent does not invest in new machines or retire old capital in response to the price of energy, which is consistent with this assumption, and suggests that the bias may not be large.

This paper estimates the strength of the relationship between the price of energy, technology and energy demand. Theoretical work on price-induced technological change began with Hicks (1932), who argued that technology is directed towards factors with high prices. Work on induced innovation (e.g. Binswanger, 1974) and directed technological change (e.g. Acemoglu, 2002) has refined the original analysis, and predicts that relative profit incentives raise the demand for certain types of technology. These models typically focus on innovation and abstract from the adoption decision. However, a finding of limited adoption would imply a limited amount of innovation.

Several recent papers study the effect of the price of energy on innovation and technology adoption.[5] This work focuses on specific technologies in a few industries, and it is difficult to extrapolate the results to the sector- or economy-wide level. Popp (2001 and 2002) examines induced innovation and technology adoption in several manufacturing industries. Popp estimates

---

[5]Largely for reasons of data availibility, early empirical tests of these models were concentrated in the agricultural sector (e.g., Griliches, 1957, and Hayami and Ruttan, 1970).

the elasticity of energy-saving patents to the price of energy, and the effect of these patents on energy efficiency. Similarly to this study, Popp finds that one-third of the observed decline in energy use was due to induced innovation. Thus, in both the energy intensive industries studied by Popp and across the entire manufacturing sector (considered in this paper), technological change has a limited role in explaining observed changes in energy demand.

Doms and Dunne (1995) and Pizer *et al.* (2002) find that high energy prices cause the adoption of specific types of energy-saving technologies in several manufacturing industries. They focus on the adoption of a variety of energy related technologies. These studies implement a different empirical strategy from this paper, exploiting cross sectional variation of energy prices in the late 1980s and early 1990s. The authors conclude that energy prices have a significant effect on technology adoption.[6]

Several studies have examined the effect of energy prices on technology adoption in other areas of the economy. Jaffe and Stavins (1995) find a positive response of the adoption of thermal insulation technology to energy prices, though the magnitude is small compared to the effect of other variables. Rose and Joskow (1990) conclude that electricity generators adopt fuel saving technology in response to an increase in the price of fuel. Newell, Jaffe and Stavins (1999) report that air conditioners and water heaters became more energy efficient after increases in the price of energy.

A few studies have investigated the correlation between plant age and technology adoption, finding mixed results.[7] Dunne (1994) reports that younger plants are not more likely to have certain advanced technologies in four 2-digit SIC industries. Luque (2000) finds that small young plants are more likely to adopt technology than small old plants. These authors do not explicitly differentiate between entrants and incumbents (e.g. Luque separates plants into 15-year age groups). Future work is needed to reconcile their findings with the results of this paper, which imply that new plants use different technology from older plants.

This paper is organized as follows. Section 2 develops a model of technology adoption and derives the estimating equation. Section 3 describes the measurement of the price of energy, and

---

[6]Also related to this paper, Anderson and Newell (2004) investigate the Department of Energy's (DOE) Industrial Assessment Centers program, in which the DOE performs audits of manufacturing plants and suggests projects that might reduce energy requirements. Anderson and Newell find that high energy prices increase the probability of undertaking a project, though other variables, such as expected implementation costs, have larger effects. The program covers a limited number of small and medium-sized plants that request the audit, so it is difficult to extend the results to the entire sector; nevertheless, they find a small, but precisely estimated, effect of energy prices.

[7]Much of the literature on technology adoption characterizes the factors that affect adoption, such as plant size and expected profitability. Recent work has studied network effects (e.g. Saloner and Shepard, 1995) and the need for complementary inputs such as skilled labor (e.g. Caselli and Coleman, 2001). This paper takes energy prices to be exogenous to these factors, and the emphasis of this paper is on the entrant/incumbent distinction, which has received little attention in the literature.

Section 4 describes the data. The main results are presented in Section 5, which also discusses the empirical support for the main assumptions. In Section 6, I use the results to calculate the effect of technology adoption on the long run response of output to an energy price shock. Section 7 concludes.

2   Model and Identification

2.1   Motivating Theory

I present a model of technology adoption that highlights the simplifying assumptions used to derive the estimating equation. Technology adoption responds to an unexpected and permanent increase in the price of energy. I derive the short run response of energy demand to the shock, where technology is constant, and the long run response, when technology changes.

2.1.1   The Baseline Model

Consider a single industry in which plants take input prices as given and choose their energy technology and capital stock when they enter. The industry begins in the steady state, in which all input prices and the demand for the industry product are constant. I characterize profit maximization for existing plants, analyze the decisions of new plants that have entered but not yet produced, and discuss the entry decision.

Plant i produces output, Y_{it}, at time t and has a production function, $F(A_i, K_i, A_i^E E_{it}, L_{it})$, where A_i is total factor productivity (TFP); K_i is the capital stock chosen by the plant; A_i^E is a state variable, representing energy technology; E_{it} is purchased energy; and L_{it} is employment. Plant i learns A_i before entering, which is drawn from the distribution function $G(A_i)$, with support $[0, \infty)$. The production function has a constant elasticity of substitution (CES):

$$Y_{it} = A_i(\alpha_K K_i^\rho + \alpha_L L_{it}^\rho + \alpha_E (A_i^E E_{it})^\rho)^{1/\rho},$$

where $-\infty < \rho < 1$, and the elasticity of substitution is $\frac{1}{1-\rho}$. This functional form yields a simple solution, and implies that the substitution elasticity does not depend on the price of energy.[8]

The maximization problem of a potential entrant at time $t = 0$ can be written as:

$$\max_{\{E_{it}, L_{it}\}_{t=0}^\infty, A_i^E, K_i} E_0\{\sum_{t=0}^\infty \frac{1}{(1+r+\phi)^t}(p_t Y_{it} - p_t^E E_{it} - p_t^L L_{it}) - H(A_i^E, K_i) - p_0^K K_i\} \qquad (1)$$

---

[8]Although the CES functional form is essential for the results below, it can be generalized somewhat, without affecting the linearity of the estimating equation. For example, if capital and energy are combined according to a CES technology to form an intermediate input (similar to Atkeson and Kehoe), the resulting empirical specification would be identical to that used below, with the inclusion of the initial capital stock as an independent variable.

$$s.t. Y_{it} = A_i(\alpha_K K_i^\rho + \alpha_L L_{it}^\rho + \alpha_E(A_i^E E_{it})^\rho)^{1/\rho},$$

where $E_0\{\cdot\}$ is the expectation operator; p_t, p_t^E, p_t^L and p_0^K are the prices of output, energy, labor and capital; r is the discount rate; ϕ is the probability the plant is exogenously destroyed; and $H(A_i^E, K_i)$ is the cost of selecting energy technology A_i^E, where the function is increasing in both arguments. Profit maximization includes three decisions. First, a potential entrant calculates its expected profits and enters if profits are non-negative. Second, the new plant selects its energy technology and capital stock as functions of expected future prices, and it pays a sunk cost of $H(A_i^E, K_i) + p_0^K K_i$. Since the production function is constant returns to scale, I set each plant's capital stock to $K_i = \overline{K}$. Finally, in each period the plant purchases labor and energy after observing current factor prices; operating costs are $p_t^E E_{it} + p_t^L L_{it}$.

Energy technology augments energy use, but because it is chosen at the same time as capital and both are fixed, it is a feature of the capital stock.[9] In other words, this is a vintage capital model, in which plants purchase capital and choose the energy characteristic of the capital stock at the time of entry. Note that this model is more flexible than the standard putty clay model (e.g., Atkeson and Kehoe), because a plant can adjust its capital-energy ratio after selecting its technology.

I now discuss the three components of profit maximization in reverse chronological order. I analyze the steady state, in which all prices are constant, and the subscript 0 denotes the initial steady state. Each period proceeds as follows. A fraction, ϕ, of plants are exogenously destroyed.[10] New plants enter and select their energy technology. All plants, new and existing, purchase factors, produce and receive revenue.

Having chosen A_i^E, a plant solves the following problem each period:

$$\max_{E_{i0}, L_{i0}} pY_{i0} - p_0^E E_{i0} - p_0^L L_{i0}$$

$$s.t. Y_{i0} = A_i(\alpha_K \overline{K}^\rho + \alpha_L L_i^\rho + \alpha_E(A_i^E E_{it})^\rho)^{1/\rho}.$$

The first order condition for E_{i0} yields the equation:

$$\ln(E_{i0}/Y_{i0}) = \frac{\rho}{1-\rho} \ln A_{i0}^E + \frac{1}{\rho-1} \ln \widetilde{p}_0^E + \frac{1}{1-\rho} \ln(\alpha_E) + \frac{\rho}{1-\rho} \ln A_i, \tag{2}$$

where $\widetilde{p}_0^E = p_0^E/p$, is the real price of energy and E_{i0}/Y_{i0} is energy efficiency. Equation (2) shows several results. First, holding the price of energy constant, a higher energy technology (A_i^E) can

---

[9] The assumption that the plant's capital stock and energy technology are fixed captures the main features of a model with adjustment costs, in which entering plants have greater flexibility and respond more elastically to a price shock.

[10] The main results are qualitatively similar if I allow for endogenous exit and a fixed fraction of plants enter each period. The assumption that ϕ is independent of A_i is for simplicity. The main results are unchanged if, for example, ϕ is a decreasing function of A_i.

increase or reduce energy demand. When $\rho < 0$ inputs are gross complements, and a plant with better energy technology uses less energy. Second, the coefficient on the price of energy is the energy price elasticity, and since $\rho < 1$, it is negative; energy demand decreases when the price of energy increases. Third, productivity, A_i, affects the level of energy demand.

Next, I examine the maximization problem of a new plant, after it has decided to enter. I assume the following cost function for selecting energy technology A_i^E:[11]

$$H(A_i^E, \overline{K}) = \frac{\overline{K}(A_i^E)^\gamma}{\gamma(r + \phi)},$$

where $\gamma > 1$. Note that the cost is directly proportional to the plant's capital stock, $\overline{K}$.

Solving the first order condition for energy technology gives:

$$\ln A_{i0}^E = \beta_1 \ln \widetilde{p}_0^E + \beta_2, \tag{3}$$

where β_1 is a constant that corresponds to the elasticity of technology to the expected price; I discuss its interpretation below. The second term, β_2, is a constant that depends on the output price, other input prices and the plant's TFP.

To analyze the effect of technology adoption on energy demand, I combine equations (2) and (3) to obtain:

$$\ln(E_{i0}/Y_{i0}) = \sigma^T \ln \widetilde{p}_0^E + \sigma^S \ln \widetilde{p}_0^E + \beta_i, \tag{4}$$

where $\sigma^T = \frac{\rho}{1-\rho} \cdot \beta_1$, and σ^T captures the effect of price-induced technology adoption (i.e. the shift of the demand curve for energy). The energy price elasticity is σ^S, which is negative; σ^S captures movement along the demand curve for energy. Finally, β_i is a plant specific constant that depends on productivity. Equation (4) is a central result because it shows the two channels by which the price of energy can affect energy demand: technology adoption (σ^T) and input substitution (σ^S).

The entry decision completes the characterization of a plant's decisions. Expected profits are increasing in productivity, A_i, so there is a cutoff productivity, $\underline{A}_i(\widetilde{p}_0^E)$, below which a potential entrant does not enter.[12] The ratio of the number of entrants to the number of plants in the

---

[11] I assume plants have the same adoption costs. In a more general case, adoption costs could be $\widehat{H}(A_i^E, c_i, K_i) = \frac{K_i(A_i^E)^{\gamma+c_i}}{\gamma+c_i}$, where a higher c_i denotes more costly adoption. In this case the elasticity of energy technology to the price of energy (β_1) depends on c_i, where a lower c_i confers a greater elasticity.

The function $H(\cdot)$ is chosen for analytical simplicity, and the exact form is not important for the main results. For example, if the cost were proportional to the capital stock raised to an exponent, the resulting estimating equation would include the capital stock as an additional righthand side variable. As discussed below, the results are unaffected by this, or similar modifications to the baseline equation.

[12] This boundary depends on other prices, but I emphasize the effect of the price of energy because the other prices remain constant in the following analysis.

industry is $\int_{\underline{A}_i(\widetilde{p}_{E0})}^{\infty} dG(A_i)$. In the steady state, this ratio is equal to the fraction of exiting plants, ϕ. Note that there is a negative relationship between the average productivity of entrants and the number of entrants.

I consider a shock to the steady state and compare the energy efficiency of entering and existing plants. The industry begins in the steady state with the price of energy equal to p_0^E. There is an unexpected, one-time and permanent increase of p^E, to p_1^E.

For simplicity, the demand for the industry's output is completely elastic at the price p. The energy demand of plants surviving from the initial steady state is:

$$\ln(E_{i1}^I/Y) = \sigma^T \ln \widetilde{p}_0^E + \sigma^S \ln \widetilde{p}_1^E + \beta_i, \tag{5}$$

where the subscript 1 indicates that the price shock has occurred and the superscript I indicates plants surviving from the initial steady state. The constant β_i is unchanged because the output price and other input prices do not change, and σ^S and σ^T do not change because the elasticities do not depend on the factor prices (a result of the CES assumption). The energy demand of new plants is given by:

$$\ln(E_{i1}^N/Y_{i1}) = \sigma^T \ln \widetilde{p}_1^E + \sigma^S \ln \widetilde{p}_1^E + \beta_i. \tag{6}$$

where the superscript N indicates new plants.

I now use equations (5) and (6) to compare the responses of entering and existing plants to the price shock:

$$\frac{\partial \ln(E_{i1}^N/Y_{i1})}{\partial \ln \widetilde{p}_1^E} - \frac{\partial \ln(E_{i1}^I/Y_{i1})}{\partial \ln \widetilde{p}_1^E} = \sigma^T. \tag{7}$$

The difference in energy demand is equal to the effect of technology adoption. This result illustrates the LeChatelier principle, that demand is more elastic in the long run when technology can adjust. The effect of the price of energy on energy demand, holding technology fixed, is σ^S. The price increase causes energy demand to fall by $\sigma^T + \sigma^S$ in the long run. Since $\mid \sigma^T + \sigma^S \mid > \mid \sigma^S \mid$, long run energy demand is more elastic.

Profits are decreasing in the price of energy, so the cutoff productivity for entry, $\underline{A}_i(\widetilde{p}_1^E)$, increases. The price of energy is positively correlated with the average productivity of entrants.

In summary, after the price shock the energy demand for new plants is given by equation (6). Equation (5) shows the energy demand of plants that survive from the initial steady state. Technology adoption causes energy demand to fall by more for entrants than incumbents, where σ^T is the difference in the responses of the two groups. Finally, there are fewer entrants each period than before the price shock and their average productivity is higher.

2.1.2 VARIABLE ELASTICITY OF SUBSTITUTION

I relax the assumption that entering plants have the same elasticity of substitution as existing plants. For simplicity, I assume that the price shock coincides with a shock to ρ, changing it to ρ_1. This setup is sufficient to show that the difference in the effect of the shock for entrants and incumbents depends on technology, as before, but also on the different substitution elasticities.[13]

The energy-price elasticity of plants remaining from before the shock is σ^S. For entering plants, the parameters from equation (4) differ from the baseline case because ρ changes. The long run elasticity is $\widetilde{\sigma}^S + \widetilde{\sigma}^T$, with $\widetilde{\sigma}^S$ and $\widetilde{\sigma}^T$ defined similarly to before, replacing ρ with ρ_1. Thus, the comparison of the elasticities at $t = 1$ is:

$$\frac{\partial \ln(E_{i1}^N / Y_{i1})}{\partial \ln \widetilde{p}_1^E} - \frac{\partial \ln(E_{i1}^I / Y_{i1})}{\partial \ln \widetilde{p}_1^E} = \widetilde{\sigma}^T + (\widetilde{\sigma}^S - \sigma^S). \tag{8}$$

There are two terms: the effect of technology adoption and the difference between the price elasticities of entrants and incumbents. In Figure 1, $\widetilde{\sigma}^T$ corresponds to the inward shift of the demand curve and $\widetilde{\sigma}^S - \sigma^S$ reflects the rotation of the demand curve. An important empirical question is whether the price elasticity is constant, i.e., whether $\widetilde{\sigma}^S - \sigma^S = 0$.

2.2 EMPIRICAL SPECIFICATION

I derive the estimating equation from the model in the previous section. I discuss the assumptions under which the comparison of entrants and incumbents yields an accurate measure of the effect of technology adoption on energy demand.

I return to the case where ρ is constant. The log energy efficiency of an entering plant at time t is:

$$\ln(E_{it}^N / Y_{it}) = \sigma^T \ln \widetilde{p}_t^E + \sigma^S \ln \widetilde{p}_t^E + \beta_i, \tag{9}$$

The current price of energy affects the energy efficiency of an entrant through the choice of technology and factor substitution.

In comparison, the log energy efficiency of an incumbent at time t is:

$$\ln(E_{it}^I / Y_{it}) = \sigma^T \ln \widetilde{p}_0^E + \sigma^S \ln \widetilde{p}_t^E + \beta_i, \tag{10}$$

where $\widetilde{p}_0^E$ is the real price of energy at the time the plant entered. The price of energy affects the energy efficiency of an existing plant via substitution; by assumption, it cannot adopt technology.

---

[13]This case yields the same conclusion as a more general model, in which plants choose technology from a menu of ρ and A^E.

I define the constant N_{it}, a dummy variable equal to one if the plant is an entrant, and combine equations (9) and (10):

$$\ln(E_{it}/Y_{it}) = \sigma^T N_{it} \ln \widetilde{p}_t^E + \sigma^S \ln \widetilde{p}_t^E + \beta_i + (1 - N_{it})\sigma^T \ln \widetilde{p}_0^E. \tag{11}$$

Thus, three factors determine energy demand: technology adoption by entrants, substitution, and a plant-specific constant, which depends on TFP and whether or not the plant is an entrant. Note that by the CES assumption, the price elasticities are equal in equations (9) and (10), allowing me to obtain equation (11).

I add an error term and a matrix of controls, and assume that the TFP component of β_i is uncorrelated with the other variables. The real price of energy varies by industry (j), state (s) and year (discussed below). The estimating equation is:

$$\ln(E_{it}/Y_{it}) = \delta_1 N_{it} \ln \widetilde{p}_{jst}^E + \delta_2 \ln \widetilde{p}_{jst}^E + \delta_3 N_{it} + X_{it}\eta + \varepsilon_{it}, \tag{12}$$

where $\{\delta_1, \delta_2, \delta_3, \eta\}$ are parameters to be estimated, X_{it} is a matrix of controls with coefficient vector η, and ε_{it} is a random disturbance term. The matrix X_{it} includes a constant, industry-region-year interactions, state fixed effects and a set of cohort dummies, equal to one if the plant entered in the corresponding year and zero in the year of entry. The cohort dummies account for the dependence of the intercept on the price of energy at the time of entry. The industry-region-year interactions control for industry-region shocks.

The parameter δ_1 measures the effect of the price of energy on the differential log energy efficiency of entrants and incumbents. Under the simplifying assumptions discussed below, this captures the effect of price-induced technology adoption on energy demand. δ_2 measures the energy price elasticity for entrants and incumbents, and the concavity of the production function implies that δ_2 is negative. δ_3 is the average difference in energy efficiency between entrants and incumbents.

There are several issues related to the identification of the effect of technology adoption on energy demand. To address the possible endogeneity of the price of energy, I compute a fixed-weight price index, described in Section 3. Industry shocks, technological change and substitution should not affect the price index.

Equation (12) includes the current price of energy. This specification allows for the simultaneous measurement of the price elasticity and the effect of technology on energy demand. This is the appropriate measure if the price of energy follows a random walk process. In the sample, the average plant operates for about 10 years, and other work (e.g., Pindyck, 1999) suggests that a random walk is a reasonable approximation over that time horizon. Mean reversion would imply

10

that the results underestimate the effect of technological change. Below I relax the random walk assumption and use a vector autoregression (VAR) to forecast energy prices, which yields similar results.

I now discuss the three simplifying assumptions. First, the price of energy is uncorrelated with the difference in the average TFP of entrants and incumbents. Several approaches discussed below suggest that the price of energy is not strongly correlated with the productivity of entrants and incumbents.

The second assumption is that entering and existing plants have an equal and constant substitution elasticity. Diamond, McFadden and Rodriguez argue that an identifying assumption, such as this one, is necessary to distinguish technological change from a change in the substitution elasticity. If the asssumption does not hold, the estimate of the shift in the demand curve would be biased for two reasons. First, as equation (8) shows, δ_1 would measure the combined effect of a rotation and shift of the energy demand curve (see Figure 1).

The other concern is that the substitution elasticity may vary along the energy demand curve (as in the case of a translog production function). Suppose entrants and incumbents have the same technology, and the price of energy does not affect technology adoption. If the price elasticity were not constant, a price increase would cause both entrants and incumbents to move along their demand curve, changing the elasticities of both groups, say from σ_0 to σ_1. The estimate of δ_2 would reflect the average of σ_0 and σ_1, while the estimate of δ_1 would equal the difference between σ_1 and δ_2, and would be statistically significant.

It appears, however, that the estimate of δ_1 reflects a shift of the demand curve. The results below support the assumption that the price elasticity of entrants and incumbents is the same, and that allowing the price elasticity to vary across time and industries does not affect the results.

The final assumption is that existing plants do not adopt technology. Technology adoption by existing plants would bias δ_1 downwards, so this assumption concerns the estimate's accuracy, and not a possible spurious correlation. It is straightforward to augment the model and allow existing plants to adopt technology after the price shock. Under reasonable assumptions, the effect of technology adoption on energy demand is greater for entrants than incumbents, i.e., $\sigma^T > \sigma^{IT}$, where σ^{IT} is the effect of technology adoption on energy demand for incumbents.[14] The total effect of technology adoption on energy demand would be equal to $\sigma^T + \sigma^{IT}$. However, δ_1 would measure $\sigma^T - \sigma^{IT}$, the difference between the effect of technology for entrants and

---

[14] For example, suppose the marginal cost of adoption is higher for incumbents, and they recover the cost $H(A_i^E, \overline{K})$ before choosing their new technology. Then A_i^E would be smaller for incumbents, and the effect of technology on energy demand would be smaller.

incumbents. This is a second illustration of the LeChatelier principle, and it is an important result because it implies that δ_1 is a lower bound of the total effect of technology adoption on energy demand.

To assess the magnitude of the downward bias, I investigate whether existing plants adopt technology. I find little evidence that investment and capital retirements respond to the price of energy, suggesting that existing plants do not adopt technology.

3 CONSTRUCTION OF THE PRICE OF ENERGY

Before discussing the data and the main results, I describe the construction and variation of the measured price of energy. I compute the nominal price of energy as the weighted sum of fuel and electricity prices:

$$p^E_{jst} = \sum_f w_{jsf} \cdot \pi_{sft}, \tag{13}$$

where p^E_{jst} is the price of energy in industry j, state s and year t, and w_{jsf} is the share of energy (in BTUs) for industry j in state s for energy source f (natural gas, electricity, coal, residual and distillate).[15] π_{fst} is the price in current dollars per million BTUs by state, source and year. I compute the real price of energy, $\widetilde{p}^E_{jst}$, by dividing p^E_{jst} by the output price for industry j in state s and year t.[16]

There is cross sectional and time series variation in p^E_{jst} because of variation in the BTU weights and the variation of the state by year energy prices. The regressions below include industry-region-year interactions and state fixed effects, so I discuss the remaining variation.

A state price shock, to one energy source or to all sources, affects relative energy prices in a given industry and region. The size of the effect varies according to the weights. On the other hand, if energy prices are uniformly high, either in a given year and region, or in one state over the entire period, this does not effect p^E_{jst}.

An aggregate or regional price shock can affect relative energy prices within a given industry-region-year cell because of the BTU weights. For example, in the Nitrogen Fertilizer industry (SIC 2873), the natural gas BTU share is lower in Ohio than in other states in the same region. In response to a regional shock to natural gas prices, the price of energy for Fertilizer plants in Ohio would fall compared to Fertilizer plants in the rest of the region.

---

[15]In 1975 the average BTU shares across states and industries were 0.56 for natural gas, 0.17 for electricity, 0.11 for coal, 0.11 for residual, and 0.05 for distillate.

[16]I also construct the price of electricity and fuels separately. A plant's real electricity price is the price of electricity in the corresponding state and year, divided by the output price. The nominal fuel price is computed similarly to equation (13), except that each weight is the BTU share of the fuel in total fuel use for the industry and state.

Fixed geographic differences in the prices of individual energy sources may cause variation in the price of energy. For example, the price of coal is relatively high in Maine. Since cement plants (SIC 3241) have a higher coal share than the average industry, the price of energy for cement plants in Maine is higher than for other plants in the same state. Cement plants in Maine would face a high price of energy relative to other cement plants in the Northeast.

It is important to note that the BTU weights, w_{jsf}, do not change over time. Thus, industry-specific substitution between energy sources and technological change, both of which are endogenous responses to energy prices, do not affect p_{jst}^E.

I briefly discuss the historical patterns of the different energy prices. Figure 2 shows the average real prices of natural gas, electricity, coal, residual and distillate from 1970-1997, with prices normalized to one in 1967. Natural gas and oil prices moved together closely, and electricity prices followed a similar pattern, but varied less over time. Coal prices show a different pattern, peaking in the early 1970s and declining afterwards.

Natural gas prices varied dramatically over time and space, due to regulation, transportation costs and oil prices. In the 1970s, federal regulation reduced the domestic supply of natural gas and contributed to a severe shortage and price increases (see MacAvoy and Pindyck, 1975).[17] Beginning in the late 1970s, prices were gradually deregulated, and by the end of the sample, prices mainly reflected differences in transportation costs.

Electricity prices were regulated over most of this period at the state level, and did not respond quickly to demand. They depended, *inter alia*, on fuel costs (most often coal), capital costs and the composition of electricity generators (e.g. coal-fired versus hydropower). There was considerable geographic, and some inter-temporal, variation in the availability of hydroelectric and nuclear power.

Coal prices varied differently from other sources due to the geographic location of mines and environmental regulation. Transportation costs are particularly high for coal, relative to extraction costs, making geographic variation an important determinant of relative prices. Sulfur emissions regulation in the 1970s and 1990s raised the demand for low sulfur coal supplied from the west. The regulation affected coal prices in a much different manner from the OPEC shocks and natural gas shortage. Most industries use little coal directly, though coal prices had a large effect on electricity prices.

Import prices of oil, transportation costs and regulation were the main sources of variation for petroleum product prices. For much of the period, OPEC held nominal crude oil prices stable, except during the two major shocks in the 1970s. Regulatory bodies such as the Texas

---

[17]During the shortage, supply was completely halted to many customers, and observed prices may not reflect the true costs. Below I find that potential mis-measurement of true natural gas prices does not affect the results.

Railroad Commission further maintained price stability by controlling the domestic supply (see Hamilton, 1984). The OPEC shocks affected the price of energy for plants in states with a high cost share of distillate or residual, relative to other states in a given industry and region. Natural gas is a close substitute for oil, which explains the high correlation of these prices across states and time.

In the baseline regression I assume that the sources of variation for all energy sources are exogenous. This assumption is supported by the importance of transportation costs and the actions of government regulators, who were unlikely to respond to changes in energy demand. I obtain similar results using the aggregate price of energy, which does not rely on this assumption.

4 DATA SOURCES AND VARIABLE CONSTRUCTION

I use several data sources: the Census of Manufactures (CM), 1963, and every five years from 1967-1997; the Annual Survey of Manufactures (ASM), 1973-1988; the ASM fuel surveys for 1975-1981; the Manufacturing Energy Consumption Survey (MECS), 1985-1998; and the Department of Energy (DOE) State Energy Price Report, 1970-2000.[18] The CM, ASM and MECS contain plant level data and are confidential.

The main variables are energy efficiency, the real price of energy and the entrant dummy variable. Energy efficiency is measured in units of real energy consumed to real output, where prices are normalized to one in 1967 (i.e., energy efficiency is equal to the cost share of energy in 1967). The CM contains output and energy expenditure in nominal dollars, so I primarily describe the price deflators. I construct a detailed product price index, which varies both geographically and by SIC product code. In each Census year there is a product file containing revenues and quantities for a large number of 7-digit products. Missing or imputed quantity data and concerns about data quality prevent me from using plant level product prices. Instead, I use the least aggregated prices available for each plant. I begin with 7-digit by state prices, and compute chain-weighted prices for different aggregation levels, up to 4-digit product code. For each product a plant produces, I use the least aggregated price for which there are at least 50 observations in the product files. A plant's shipments deflator is the weighted sum of the product prices.[19]

In comparison, most other studies use the 4-digit industry prices from the NBER Manu-

---

[18] The DOE prices are available for all sources from 1970-2000. The DOE has publishted pre-1970 data for some sources, which does not appear to be comparable. I assume that nominal energy prices were constant between 1967-1970, which is supported by the trends in the pre-1970 prices.

[19] In computing the prices I drop prices below the 5th and above the 95th percentiles in an industry-state-year cell to limit the effect of outliers. The main results are insensitive to changing either the cutoff percentile or the number of plants required to use a price.

facturing Productivity Database (MP).[20] I construct the more detailed product prices for three reasons. First, there is considerable geographic variation of output prices in the data, both within and across years. Second, because I use state-level energy price data, the detailed prices prevent all geographic variation in the real price of energy from being attributed to the DOE prices and energy weights. Third, I deflate each plant's products by an appropriate product deflator, rather than an industry deflator. This approach yields a more accurate price deflator for plants producing a large share of products from other industries. The correlation between the product prices and the NBER prices is high. As the results show below, using the NBER prices instead of the CM- and DOE-based prices yields similar estimated magnitudes, though the precision is considerably lower.

To construct the dependent variable, log energy efficiency, I compute an ideal price index for energy. I use the DOE state prices and state by industry energy input cost shares, computed from the ASM and MECS for 1975-1997. I assume the shares were constant before 1975, and I extrapolate between years when there is no survey available. The dependent variable is the log of the ratio of real energy use (energy expenditure divided by the energy price index) to real output (reported shipments divided by the shipments deflator).

The previous section described the construction of the log real price of energy, which is an independent variable. The nominal price of energy is the weighted sum of the DOE state energy prices (see equation (13)). The weights are the average BTU shares computed from the 1975 ASM fuel survey, which vary by industry and state.[21] I deflate the nominal energy price by the industry by state output price. The latter is the weighted sum of product prices across an industry and state.

In summary, the dependent variable and independent variable use different deflators. The dependent variable uses plant specific product prices and an ideal price index for energy, computed from cost shares that vary over time. The independent variable uses industry by state output prices and a fixed weight energy price index, using BTU shares.

The entrant dummy variable is equal to one when a plant appears in the CM for the first time. I construct the variable using the plant identification number, which links observations of each plant over time.[22] As a result, I measure entry for some plants several years after they

---

[20] An exception is Foster *et al.* (2004), who find that geographic variation in output prices greatly affects estimated productivity.

[21] The fixed weights imply that the measured price of energy over-estimates increases in the true price index. This would bias the results if the measurement error were correlated with the entrant dummy. This does not appear to be a significant concern, however. The price of electricity and electricity efficiency do not depend on these weights, and the results are similar for electricity and total energy (see Table 7).

[22] There may be some measurement error for this variable, as plants can appear in the CM that do not operate. Using employment to impute entry (following Davis, Haltiwanger and Schuh, 1996) yields almost identical results.

enter, and do not include plants that enter and exit between Census years. Data limitations prevent me from measuring entry more precisely for the entire sample.

The sample includes all plant-year observations in the CM, from 1967-1997 (I use the 1963 Census to identify entrants in 1967). I drop "administrative record" plants, as is customary in using the CM. These are typically small plants for which most variables except payroll have been imputed. I also omit observations with non-positive output or energy expenditure. The dataset includes 1,284,597 observations for 563,961 plants.

Table 1 provides key descriptive statistics for the CM data. For each Census from 1967-1997, I separate plants according to whether they entered that year. The table reports the mean and standard deviation of log energy efficiency, unweighted and weighted by plant output, and the nominal cost share of energy (energy expenditure divided by total shipments).[23] Recall that prices are normalized to one in 1967, so that each plant's cost share is equal to its energy efficiency that year. The different measures show a common pattern: the difference in energy efficiency between entrants and incumbents increases when the price of energy is high.[24]

Figure 3 shows the same pattern graphically, for 1972-1987.[25] The series for incumbents is the output-weighted mean log energy efficiency of plants that operated in the previous Census (plants for which the entrant dummy in equation (12) is zero). Incumbents in 1972 include plants that entered in 1967 and plants still operating from the 1963 census. The other lines show the mean energy efficiency by entry cohort.

The figure shows that entering plants are relatively more efficient than incumbents when the price of energy is high. The 1972 cohort, which entered before the price shock, is 0.17 more efficient than incumbents in 1972. The average log energy efficiency of the 1977 cohort, most of which entered before or during the first shock, is about 0.29 below the incumbents, but is similar to the 1972 cohort. In contrast, the 1982 cohort is further below the incumbents and other entry cohorts. The 1987 cohort, which entered as energy prices fell, is closer to the incumbents than is the 1982 cohort. The 1982 cohort remains more efficient than the other groups in 1987. As discussed above, one concern is that entrants may have a different substitution elasticity than incumbents. The fact that the 1982 cohort does not increase energy demand by more than the other groups suggests that this is not the case.

---

[23]There is a considerable amount of variation across plants in the same entry cohort and year, some of which appears to be due to reporting errors. In particular, the standard deviation of the cost share is often quite large. This is driven by the presence of a few outliers; for example, in no year do more than 100 plants report a cost share larger than 1. As shown in Table 2, omitting these outlying plants from the regressions does not affect the results.

[24]The median energy efficiency (not reported) shows a similar pattern.

[25]I omit the other years for clarity.

5.1   BASIC SPECIFICATION

This section discusses the estimated coefficients from equation (12), using plant data from 1967-1997. The model is estimated by Weighted Least Squares, using the plant's share in shipments for the corresponding year to account for the reporting errors of small plants.[26] Standard errors are robust to heteroskedasticity using the Huber-White formula. Column 1 of Table 2 shows the baseline specification, which includes industry-region-year and state dummies to control for demand shocks and fixed differences across states. The dependent variable is log energy efficiency and I include a full set of entry cohort dummies. The estimate of the interaction of the log price of energy and the entrant dummy is -0.100, with standard error 0.026 (significant at the one percent level), meaning that the relative efficiency of entrants improves when the price of energy rises. The estimate of δ_2, the energy price elasticity, is -0.172 with standard error 0.019.[27] The price of energy is demeaned in all regressions, and entrants are about 8 percent more efficient on average.

The magnitudes imply that technology explains a significant, though relatively small amount of the observed changes in energy demand. For example, between 1972 and 1982 the real price of energy rose 150 percent while energy efficiency improved by 33 percent. The estimate of the interaction coefficient implies that about 30 percent of the observed change in efficiency was due to technology adoption. These results agree with Popp (2001), who performs an analogous calculation for a set of energy intensive industries, and finds that one-third of the observed decline in energy use was due to innovation.[28] In Section 6 I return to the question of how much technology adoption reduces the long run effect of a price shock on GDP.

Recall that some cross sectional variation of the price of energy is due to fixed differences in industry weights and energy prices. In column 2, I include a full set of industry-state interactions, to eliminate cross sectional variation in energy prices. The estimate of δ_1, -0.077, is similar to the baseline, and is significant at the 10 percent level. This suggests that transportation costs and other time-invariant factors are not driving the results.

It is possible that the specification in column 1 does not adequately account for industry heterogeneity. In column 3, I add to the baseline specification the interactions of the log price

---

[26] Dhrymes (1992) finds that there is less measurement error for large plants. Presumably this is because the Census Bureau focuses on publishing aggregate statistics. It checks the responses from large plants more carefully for missing values or nonsensical answers.

[27] It appears that this estimate reflects a plant's ability to adjust its energy demand, rather than cross-plant heterogeneity. The price elasticity measured from a first differenced regression of energy efficiency on the price of energy (including industry-region-year interactions) is similar to the estimate of δ_2 in Table 2.

[28] Popp's estimate of δ_1 is much larger, -0.37, but because the industries in his sample experience a larger improvement in energy efficiency, he obtains a similar estimate of the importance of technological change.

of energy with a set of industry dummies. This allows each industry to have a different average log efficiency in each region and year, and a different price elasticity over the entire period. The estimate of δ_1 is quite similar to the baseline.

Next, I relax the assumption that plants in the same entry cohort have the same average log energy efficiency. In column 4 I interact the cohort dummies with a set of 2-digit industry dummies, which has little effect on the estimates.

Finally, since the model is estimated by Weighted Least Squares, it is possible that outlying observations affect the results. Reporting errors are common in the CM, and are a particularly likely source of measurement error. To address this concern I calculate the 5th and 95th percentile of the energy cost share (energy expenditure divided by value of shipments) for each industry-region-year-entrant cell. Column 5 omits outlying observations, and the results remain robust.[29]

The magnitude of the energy price-entrant interaction, δ_1, is relatively stable and statistically significant after accounting for industry, geographic, and entry-cohort heterogeneity, and does not appear to be affected by outliers. The remaining specifications assess the validity of the simplifying assumptions used to derive equation (12).

5.2 SAMPLE SELECTION AND PLANT HETEROGENEITY

The assumption that the substitution elasticity is equal for entrants and incumbents is central to the identification of the baseline equation. I postpone that discussion to the next section, because it is first necessary to assess whether the sample used in Table 2 is representative.

Since the productivity distribution depends on entry and exit, there is a problem of sample selection. Existing approaches, such as a Heckman correction or methods based on the propensity score (e.g., Angrist, 1995) require a sample where the selection decision is observed and the independent variables are uncorrelated with the unobserved variables. The CM does not meet these conditions, because I do not observe potential entrants that do not enter.

I consider a regional market for entry, analogous to Dumais, Ellison and Glaeser (2002). I assume that the distribution from which potential entrants are drawn, $G(A_i)$, is constant in a given industry, region and year. In the model in section 2.1, the fraction of entrants at time t is $\int_{\underline{A}_i(p_t^E)}^{\infty} dG(A_i)$. If the cutoff productivity, $\underline{A}_i(p_t^E)$, is increasing in the price of energy, a high price of energy would cause the entry of fewer, more productive plants. The average entrant would be more efficient when the price is high, introducing the possibility of a spurious correlation

I use the log of a count of entrants for the plant's state, industry and year of entry (the

---

[29] The specification in column 5 omits outliers due to reporting errors, since the energy cost share is computed directly from values in the CM. Alternatively, the results may be affected by outlying values of energy efficiency; dropping these observations in a similar manner has little effect on the results.

same level of variation as the price of energy) to proxy for unobserved productivity. A negative correlation of this variable with the price of energy would suggest that a high price of energy causes the entry of more productive plants.

I construct a second test using the survival of entrants. Given a stable distribution of potential entrants, plants from an entry cohort with greater than average productivity would be more likely to survive to the next Census. A positive correlation of the price of energy with survival would indicate that the price of energy is positively correlated with the productivity of entrants.

5.2.1 Entry and Survival of Entrants

Table 3 shows the correlations between the log price of energy and entry, survival of entrants and exit. In all regressions, the dependent variable is aggregated to the industry-state-year level. I include state fixed effects and a full set of industry-region-year interactions. I weight observations by the total shipments of plants in the corresponding cell. In column 1, the dependent variable is the log of the count of entrants, and the coefficient on the log price of energy is 0.059, with standard error 0.027. As noted above, a negative coefficient would suggest a spurious correlation in Table 2, so this does not appear to be a significant concern.

The positive correlation in column 1 implies that the results may be biased downward. However, there are two concerns with that regression. First, there are a large number of cells with zero entrants, which are not included in the log specification. Second, energy prices may be positively correlated with the number of plants operating in a state and industry. Industry-state-year cells with more plants may have more turnover, which would result in a positive correlation between the price of energy and entry. In column 2 the dependent variable is the ratio of entrants to the number of plants operating in the cell, which addresses both issues. The estimate is negative and insignificant, and the implied elasticity is quite small, less than -0.01. Although the relationship between entry and the price of energy is sensitive to alternative specifications, there is no evidence of a significant negative correlation.

In column 3, the dependent variable is the log of survival, and in column 4, it is survival divided by the number of entrants. I do not include industry-state cells for 1997 because I do not observe whether 1997 entrants survive. In both cases the price of energy has an insignificant effect on survival, which provides further support for the assumption that plant productivity is not strongly correlated with the price of energy.

For completeness, Table 3 includes results for plant exit. Although there is no general relationship between exit and plant productivity, large or significant coefficients might indicate bias. The dependent variable in column 5 is the log of the number of plants that exited between

the previous and current Census, and in column 6 it is the ratio of exiting plants to the initial number of plants operating. In both regressions the coefficient on the price of energy is positive, though the elasticity is much smaller and less precisely estimated using the exit fraction. If less productive plants are more likely to exit, the estimate in column 5 would imply that the baseline results are biased towards zero. In the sample, there is a weak correlation between the number of plants exiting and the average energy efficiency of those plants, which implies that this bias is likely to be small.

5.2.2 Controls for Unobserved Heterogeneity

Table 4 reports several regressions that include proxies for unobserved heterogeneity. Columns 1 through 4 add as controls to equation (12) the dependent variables from columns 1 through 4 in Table 3. I use the entry or survival measure in the year the plant entered; the variable is constant during the life of the plant. These variables proxy for the average productivity of entrants in the industry year and state the plant entered. In each case the estimate of δ_1 is unchanged from the baseline result.[30] The coefficient on log entry in column 1 is small and marginally significant, which suggests that this variable is not strongly correlated with the average efficiency of entrants, as would be the case if sample selection were driving the results. The survival variables in columns 3 and 4 are negative and significant, in agreement with the hypothesis that plants with greater TFP use less energy and are more likely to survive, but this relationship does not affect the main estimates.

Recent methods of controlling for unobserved heterogeneity rely on predicted relationships between unobserved and observed variables. For example, Olley and Pakes (1996) argue that log investment can proxy for unobserved productivity, but in this sample missing values (i.e., zero reported investment) would greatly reduce the sample size. Levinsohn and Petrin (2001) use intermediate materials, but the industries in the sample do not support the separability assumption that this method requires. Consequently, I do not report the results from implementing these methods, though the baseline results are unaffected by including the log of investment or materials.

Other research in the manufacturing sector (e.g. Audretsch and Mahmood, 1995) has shown that plants within the same year and industry enter with different capital stocks and employment, possibly due to productivity differences. In columns 5 and 6, I control for the log capital stock and log employment the year the plant enters. The regression in column 5 spans 1972-1992 because

---

[30]The baseline estimates for these samples are almost identical to the estimates in columns 1-4. For example, the estimate of the interaction term using the baseline specification for the sample in column 4 is -0.099 with standard error 0.017.

of data availability.[31] Both of the productivity proxies are positive and precisely estimated, but the main coefficients of interest are unaffected.[32]

5.3 SUBSTITUTION ELASTICITIES OF ENTRANTS AND INCUMBENTS

The CES assumption implies that entrants and incumbents have the same substitution elasticity, and that the substitution elasticity is constant for a given technology. Table 5 shows that relaxing the CES assumption does not affect the results.

Section 2.1.2 shows that if entrants' technology allows for easier substitution, δ_1 would overestimate the shift of the demand curve. In column 1 of Table 5, I test whether entrants in a given year have a different price elasticity from incumbents in the corresponding year. I estimate the following equation:

$$\ln(E_{it}/Y_{it}) = \phi_0 \ln p_{jst}^E + \sum_t (\phi_t O_{it} \ln p_{jst}^E + \varphi_t T_{it} \ln p_{jst}^E) + \sum_t (\iota_t O_{it} + \kappa_t T_{it}) + X_{it}\eta + \varepsilon_{it}, \quad (14)$$

where O_{it} is a set of dummy variables, equal to one if the plant operates in the corresponding year. Thus, O_{it} is constant over the life of the plant, and a plant that operates from 1972-1977 will have $O_{i72} = 1$ and $O_{i77} = 1$. T_{it} is a set of dummy variables equal to one if the plant enters that year, and is also constant over the life of the plant. The term $O_{it} \ln p_{jst}^E$ is the interaction of O_{it} with the price of energy, and similarly for $T_{it} \ln p_{jst}^E$. This specification estimates a separate price elasticity for each year, maintaining the assumption that existing plants do not adopt technology. The energy price-entrant interaction allows the price elasticity of an entry cohort to differ from that of the corresponding population of incumbents. Note that because O_{it} and T_{it} are constant over the life of a plant, this regression does not compare the energy efficiency of entrants with all incumbents, as in the baseline specification. Rather, it compares the price elasticities of different sets of plants after they have entered, i.e., their short run price elasticities.

In column 1 I include observations from 1972-1997 because the log real energy price is normalized to zero in 1967. The price-entrant interactions are insignificant and small in magnitude, suggesting that entering plants have a similar price elasticity to incumbents. In addition, the estimate of δ_1 (not reported) is unaffected if I estimate the baseline estimating equation, (12), omitting plants that operated in a given year (e.g., all plants for which $O_{i72} = 1$). These results suggest that the price elasticities of entrants and incumbents are similar, and that the estimate of δ_1 is unbiased.

---

[31] More specifically, the capital stock for entrants is only available for those years. I set the initial capital stock of a 1972 incumbent equal to its capital stock in 1972.

[32] The estimate of the interaction terms using the corresponding samples with the baseline specification are -0.101 (standard error 0.019) and -0.092 (standard error 0.02).

A second concern is that the substitution elasticity may not be constant for a given technology. An increase in the price of energy would cause the elasticity of existing plants to change as they move along their demand curve. Entrants would have a different elasticity from incumbents on average, even if they use the same technology. Columns 2-4 return to the baseline specification, using observations from 1972-1997, and allow the price elasticity of incumbents to vary over time. Since energy prices varied significantly from 1972-1997 (see Figure 2), the estimate of δ_1 would change if the price elasticity varies with the price of energy. Column 2 includes price of energy by year interactions, which allows for a different average substitution elasticity for incumbents each year. The specifications in columns 3 and 4 estimate a separate price elasticity by region and year, and by 2 digit industry and year. In each case the estimate of δ_1 is quite similar to the baseline.

5.4 ROBUSTNESS

5.4.1 ALTERNATIVE MEASURES OF THE PRICE OF ENERGY

In Table 6, I use several alternative measures of the price of energy and find similar results. As discussed earlier, much of the variation in the price of energy is driven by aggregate shocks to the price of oil and the natural gas shortage. The geographic variation reflects regulatory and transportation differences, for which the exogeneity argument may be weaker. In column 1, I use the log aggregate price of energy, which is the log of a weighted average of the industry by state energy prices. The estimate of the interaction is -0.045 with standard error 0.012, which is smaller than the estimate in column 1 of Table 2. The fact that the estimate is negative and significant suggests that the main results do not arise purely from cross sectional price variation.[33]

Recall that in the model I consider a one-time and permanent increase in the price of energy. In that case, or if the price of energy follows a random walk, the current price is the appropriate measure of the expected future price. On the other hand, if the actual price were mean-reverting, the results would underestimate the effect of technology adoption on energy demand. In column 2 I relax the random walk assumption. I estimate a vector autoregression with the prices of natural gas, electricity, coal, residual and distillate, using two lags and state prices from 1970-2000. For a given year and energy price, the expected price is the discounted sum of the linear forecasts for the following 10 years. I weight the forecasted prices using sector-average BTU weights to obtain the aggregate forecasted price. The estimate on the interaction term in column 2 is -0.122 with

---

[33]The standard errors are similar if I use the aggregate price and cluster the standard errors by year. The standard error on the interaction term is 0.013

standard error 0.020. A comparison with column 1 indicates that the baseline specification may underestimate the long run effect of a price shock on energy demand, though the estimate in column 2 is similar to the baseline.[**Either use level VAR and i-s-y or mention variation**]

In columns 3 and 4, I account for the fact that some entrants operated before the corresponding Census year (e.g., 1977 entrants may have first operated between 1973-1977), and that technology adoption may respond gradually to the price of energy. In column 3, I use the five year lag price, and in column 4 the average price over the previous five years, for the plant's state and industry. Both estimates of δ_1 are close to the baseline.

Finally, in column 5 I use the energy and shipments prices from the Manufacturing Productivity Database (MP), to check the quality of the constructed prices. The MP prices vary by industry and year, and the industry-region-year interactions absorb the main effect of the price of energy. The coefficient on the interaction, -0.065, is smaller than the estimate in Table 2, and the precision is lower, supporting the use of the DOE- and LRD-based prices.

5.4.2 ADDITIONAL ROBUSTNESS CHECKS

Several studies report an asymmetric response to the price of energy. For example, Davis and Haltiwanger (2001) find that employment responds more to a price increase than to a decrease. Learning by doing may explain an asymmetric effect of technology adoption.[34] To investigate this possibility, I compute a dummy variable equal to one if the price of energy is at least as large as it was in the corresponding state and industry, five years earlier. The interaction of this variable with the other independent variables allows for both an asymmetric price elasticity and an asymmetric effect of technology adoption on energy demand. I find little evidence for an asymmetric effect of technology. The estimate of the average effect of technology adoption is -0.101, with standard error 0.037 (unreported). The magnitude of the additional effect of a price increase is small (0.014) and insignificant.

Table 7 reports several additional results. The standard errors in Table 2 do not account for correlation of observations by year or plant. Columns 1 and 2 show that clustering the standard errors by year or plant does not affect the precision of the estimates.

Fuel-saving and electricity-saving technology adoption may respond differently to prices, because of different available technologies and different amounts of price variation (see Figure 2). In columns 3 and 4, I separate energy use into electricity and fuel consumption, divided by

---

[34]To see this, consider an initial period of low energy prices, when the cost of adopting energy technology is relatively high. If the price increases and plants adopt better technology, the cost of future adoption decreases. If the price subsequently falls to its initial level, plants may use better technology than before the initial shock because adoption is less costly.

output.[35] The results for fuel and electricity efficiency are similar to the results for total energy efficiency.

Other work with the Census of Manufactures (e.g., Dunne, Roberts and Samuelson, 1989) suggests that plants belonging to multi-plant firms perform differently from single-plant firms. For example, entering plants that are part of pre-existing firms may have a better knowledge of energy-saving technology. To address this possibility, I construct a dummy variable equal to one if the entering plant is owned by a new firm. I interact the dummy with the log price of energy and include both variables in the baseline specification in column 5. The entrant dummy-price interaction measures the effect for all new plants, and the new firm dummy-price interaction estimates the differential effect for plants that belong to new firms. The coefficient on the energy price-new firm interaction is negative but insignificant (not reported) and the estimate of δ_1 is similar to the baseline.

In column 6, I consider whether the effect of technology adoption varies by an industry's energy intensity. I omit industries in the top quartile of the distribution of energy cost share. The estimates are robust to restricting the sample in this manner.[36]

Beginning in the early 1980s, many plants produced electricity and sold the surplus. For these plants, changes in the price of energy affected output as well as input prices. I do not have data on electricity sales, so in column 7 I drop the four 2-digit industries for which electricity sales comprise the largest share of output, as identified in the Edison Electric Institute publication, Capacity and Generation of Non-Utility Sources of Energy (1991). These industries are Paper, Chemicals, Petroleum and Primary Metals. The estimates are similar to the baseline.

In column 8, I exploit the fact that the variation of coal prices was different from other energy prices. This provides a test for the existence of an omitted variable correlated with the price of energy. Consider an omitted variable correlated with natural gas and oil prices, which affects the energy efficiency of entrants. Given the relatively low correlation between coal and oil prices, it is unlikely that this variable would be strongly correlated with coal prices; the omitted variable bias would be smaller for plants in coal intensive industries. In column 8 I allow coal intensive industries to have a different response to the price of energy. I interact the main independent variables with a dummy variable, equal to one if the plant's industry is in the 90th percentile of the cost share of coal in total energy expenditure (computed from the 1975 ASM fuel file). The estimate of δ_1 is close to the baseline, and none of the coal industry

---

[35] For comparison, the estimates of δ_1 for the baseline specification, with the same samples as in columns 3 and 4, are -0.104 and -0.105.

[36] During the 1970's, the Federal government subsidized research and development aimed at reducing energy demand. This funding was aimed exclusively at energy intensive industries, so the specification in column 6 is a joint test of heterogeneity across industries and the effects of the subsidy.

interactions are significantly different from zero, which suggests that these industries responded similarly to the price of energy.

During the natural gas shortage, supply to many areas was eliminated, and the observed price of natural gas may not always reflect the true cost. Since I cannot observe the availability of natural gas to an individual plant, I use the ASM fuel survey to identify counties in 1975 with low natural gas use. In column 9, I omit observations of plants in counties belonging to the bottom 15th percentile of the distribution, and the estimates are unchanged from the baseline.

Finally, plants may substitute towards intermediate materials when energy prices rise. For example, plants may find it less expensive to import some energy intensive goods, rather than producing the materials themselves. Since the dependent variable uses gross output, this substitution would appear to increase energy efficiency, even though output per worldwide energy consumption has not changed. This does not appear to be a major concern, as shown in column 10; estimating the baseline specification with value added (gross output net of materials and energy expenditure) yields identical results.

5.5 Technology Adoption by Existing Plants

If entering plants adopt technology, it is likely that they would report changes in capital investment or retirements. I construct a balanced panel of plants in the CM and ASM from 1972-1988. I find little evidence of technology adoption by existing plants.

The balanced panel consists of 12,014 plants. I compute each plant's capital stock using the book value of machines in 1972 and the perpetual inventory method. Log investment is the first difference of the log capital stock to account for plant-year observations with zero investment.

I cannot observe energy-related capital investment. As a proxy, I interact the price of energy with the energy cost share of a plant's industry. A positive coefficient would imply that if the price increases, investment in energy intensive industries increases, which is likely to include energy-related technology.

In Table 8 all variables are in first differences, to remove plant fixed effects, and all regressions include industry-region-year interactions. I use the current price of energy in column 1, the 5-year lag price of energy in column 2 and the aggregate forecasted price of energy in column 3, interacting the prices with the energy cost share for the corresponding industry. The dependent variable is log investment. Plants in energy intensive industries invest more when the price was high five years earlier (column 2). However, plants do not respond to the forecasted price, and there is limited evidence of technology adoption.

If a plant adopts technology, it likely retires capital that used the old technology. Retirement

data may have less measurement error than investment, and I report results using data on capital retirements for the years 1977-1988 (due to data availability). As columns 4-6 show, log retirements do not respond positively to any of the measures of the price of energy. There is little evidence that existing plants adopt technology, but without more detailed capital data, it is impossible to make further progress on this question.

6 Implications for Energy Prices and Growth

The Introduction noted the connection between the price of energy and economic activity, citing recent concerns about the prospects of rising energy prices. In this section I use the empirical estimates to calculate the effect of technology adoption on the long run relationship between GDP and the price of energy. The results suggest a limited effect of technology adoption.

I consider a 10 percent permanent increase in the price of energy, which represents the effect of declining natural resource supplies or a carbon tax. I compare the effect on steady state output for two models: one without technological change, and one with technological change.

I use the baseline model in section 2.1.1, except that I fix the number of plants in the industry. A constant number of plants enter and exit the industry in each period, and industry output is determined by output per plant.

I define the variable ε_Y^{NT} as the long run elasticity of output to the price of energy in the model without technology adoption. An increase in the price of energy causes energy consumption to decrease, lowering output. The magnitude of this effect depends on the price elasticity, and I use the estimated value of σ^S, setting σ^T equal to zero.

Analogously, ε_Y^T is the price elasticity of output for the model with technology adoption. In response to the shock, entering plants use better energy technology, which causes the demand curve for energy to shift towards the origin. This increases the level of output for a given amount of energy use, and allows output to recover in the long run. The long run effect of a price shock on output depends on the estimated values of σ^S and σ^T.

To compare the two models, I calculate the ratio $\varepsilon_Y^T/\varepsilon_Y^{NT}$, which measures the amount that long run output falls in the model with technological change, relative to model without technological change. This ratio is equal to $1 - \varepsilon_{A^E}$, where ε_{A^E}, is the elasticity of energy technology to the price of energy. This elasticity is unobserved, but can be calculated from the empirical results, using the fact that the estimate of σ^T is equal to $\frac{\rho}{1-\rho}\varepsilon_{A^E}$ (see equation (4)). Thus, $\varepsilon_{A^E} \simeq 0.12$, and the long run decrease of output in the model with technological change is about 88 percent of the decrease of the model without technological change.

In contrast, Atkeson and Kehoe simulate a putty-clay model and find that technology adop-

tion causes energy prices to have a small long run effect on output. They conclude that their results are similar to earlier empirical estimates of the long run response of energy demand (e.g., Griffin and Gregory, 1976), which use aggregate data to identify the effect of technology. These estimates may be biased upward by omitted variables, however, as the results in this paper suggest a much smaller effect of technological change. Thus, the Atkeson and Kehoe analysis may overstate the effects of technological change, though future work is needed to reconcile the results.

7 Conclusions

This study measures the effect of price-induced technology adoption on energy demand. The results have important implications for understanding the long run relationship between the price of energy and growth. A 1 percent increase in the price of energy leads to a 0.1 percent increase in the relative efficiency of entering plants versus incumbents. I find empirical support for the simplifying assumptions used to derive the estimating equation. The estimates are robust to a variety of specification checks. The calculation in Section 6 suggests that technology adoption has a small effect on the sensitivity of steady state output to the price of energy.

The results are important for two other recent discussions. First, several papers (e.g., Goulder and Schneider, 1999) have incorporated price-induced technological change into climate change models. Popp (2004) calibrates a long run growth model and calculates the net benefit of implementing a carbon tax to reduce greenhouse gas emissions. However, to accurately assess the role of price-induced technological change, these models should be calibrated to empirical estimates of the average effect of technology adoption; future work may incorporate the results reported here.

Second, recent energy price increases have renewed interest in the effects of energy prices on the value of installed capital. Alpanda and Peralta-Alva (2004) argue that the 1973/4 oil shock caused the obsolescence of energy intensive capital, which explains the coinciding decline in the stock market. In contrast, Wei (2003) argues that the effect of the price of energy on the value of capital should be roughly proportional to its cost share, and that the oil shock did not have much effect on the stock market. The results from this paper cannot be directly integrated into these models, so further work is needed to assess whether the oil shocks led to the replacement of a large fraction of the capital stock, and whether this mechanism can explain the apparently large effects of energy prices on the economy.

8 References

Acemoglu, Daron, 2002. Directed Technological Change. *Review of Economic Studies*, 69, 781-809.

Adelman, M. A., 1995. *The Genie out of the Bottle: World Oil Since 1970.* MIT Press, Cambridge, MA.

Alpanda, Sami and Adran Peralta-Alva, 2004. "Oil Crisis, Energy-Saving Technological Change and the Stock Market Crash of 1973-74." mimeo.

Anderson, Soren T. and Richard G. Newell, 2004. Information Programs for Technology Adoption: the Case of Energy-Efficiency Audits. *Resource and Energy Economics*, 26, 27-50.

Angrist, Joshua, 1995. Conditioning on the Probability of Selection to Control for Selection Bias. *NBER Technical Working Paper* #181.

Atkeson, Andrew and Patrick J. Kehoe, 1999. Models of Energy Use: Putty-Putty versus Putty-Clay. *American Economic Review*, 89, 1028-43.

Audretsch, David and T. Mahmood, 1995. New Firm Survival: New Results Using a Hazard Function. *Review of Economics and Statistics*, 77, 97-103.

Binswanger, Hans, 1974. A Microeconomic Approach to Induced Innovation. *Economic Journal*, 84, 940-958.

Capone, Charles and Kenneth Elzinga. Technology and Energy Use Before, During and After OPEC: The U.S. Portland Cement Industry. *The Energy Journal*, 93.

Caselli, F. and W. Coleman, 2001. Cross-country Technology Diffusion: The Case of Computers. *American Economic Review*, 91, 328-35.

Davis, Steven and John Haltiwanger, 2001. Sector- Job Creation and Destruction Responses to Oil Price Changes. *Journal of Monetary Economics*, 48, 465-512.

Davis, Steven, John Haltiwanger and Scott Schuh, 1996. *Job Creation and Job Destruction.* MIT Press, Cambridge, MA.

Dhrymes, Phoebus, 1992.

Diamond, McFadden and Rodriguez, 1978. Measurement of the Elasticity of Substitution and Bias of Technical Change, in *Production Economics: a Dual Approach to Theory and Applications*, vol 2, 125-148. Fuss, Melvyn and Daniel McFadden, eds., Amsterdam: North-Holland.

Doms, Mark E. and Timothy Dunne, 1995. Energy Intensity, Electricity Consumption, and Advanced Manufacturing Technology Usage. *Technological Forecasting and Social Change*, 49, 297-310.

Dumais, Guy, Glenn Ellison and Edward Glaeser, 2002. Geographic Concentration as a Dynamic Process. *Review of Economics and Statistics*, 84, 193-204.

Dunne, Timothy, 1994. Plant Age and Technology Use in U.S. Manufacturing Industries. *RAND Journal of Economics*, 25, 488-99.

Dunne, Timothy, Mark Roberts and Robert Samuelson, 1989. The Growth and Failure of U.S. Manufacturing Plants. *Quarterly Journal of Economics*, 104, 671-98.

Foster, Lucia, John Haltiwanger and Chad Syverson, 2004. Reallocation, Firm Turnover and Efficiency: Selection on Productivity or Profitability? mimeo

Goulder, Lawrence H. and Stephen H. Schneider, 1999. Induced Technological Change and the Attractiveness of CO_2 Abatement Policies. *Resource and Energy Economics*, 21, 211-53.

Greenstone, Michael, 2002. The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures. *Journal of Political Economy*, 110.

Griffin, James M. and Paul R. Gregory, 1976. An Intercountry Translog Model of Energy Substitution Responses. *American Economic Review*, 66, 845-57.

Griliches, Zvi, 1957. Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica*, 25, 501-22.

Hamilton, James, 1984. Oil and the Macroeconomy Since World War II. *Journal of Political Economy*, 91, 228-48.

Hayami, Yujiro and Vernon Ruttan, 1970. Factor Prices and Technical Change in Agricultural Development: The U.S. and Japan, 1880-1960. *Journal of Political Economy*, 77, 1115-1141.

Hicks, John, 1932. *The Theory of Wages.* London: McMillan.

Jaffe, Adam and Robert Stavins, 1995. Dynamic Incentives of Environmental Regulations: the Effects of Alternative Policy Instruments on Technology Diffusion. Journal of Environmental Economics and Management, 29: S43-63.

Levinsohn, James and Amil Petrin, 2001. Estimating Production Functions Using Inputs to Control for Unobservables. mimeo

Luque, Adela, 2000. An Option Value Approach to Technology in U.S. Manufacturing: Evidence from Plant Level Data. *Center for Economic Studies working paper #00-12.*

MacAvoy, Paul and Robert Pindyck, 1975. *The Economics of the Natural Gas Shortage, 1960-1980.* American Elsevier Publishing Company, New York.

Mansfield, E., 1968. *Industrial Research and Technological Innovation.* Norton, New York.

Newell, Richard, Adam Jaffe and Robert Stavins, 1999. The Induced Innovation Hypothesis and Energy-Saving Technological Change. *Quarterly Journal of Economics*, 114, 907-940.

Olley, Steven and Ariel Pakes. The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, 64, 1263-98.

Pindyck, Robert S., 1999. The Long Run Evolution of Energy Prices. *The Energy Journal*, 20, 1-27.

Pizer, William A., Winston Harrington, Raymond J. Kopp, Richard D. Morgenstern and Jhih-Shyang Shih, 2002. Technology Adoption and Aggregate Energy Efficiency. *Resources for the Future Discussion Paper*, 02-52.

Popp, David, 2001. The Effect of New Technology on Energy Consumption. *Resource and Energy Economics*, 23, 215-39.

Popp, David, 2002. Induced Innovation and Energy Prices. *American Economic Review*, 92, 160-80.

Popp, David, 2004. ENTICE: Endogenous Technological Change in the DICE Model of Global Warming. *Journal of Environmental Economics and Management*, 48, 742-68.

Rose, Nancy and Paul Joskow, 1990. The Diffusion of New Technology: Evidence from the Electric Utility Industry. *RAND Journal of Economics*, 21, 354-373.

Rotemberg, Julio and Michael Woodford, 1996. Imperfect Competition and the Effects of Energy Price Increases on Economic Activity. *Journal of Money, Credit and Banking*, 28, 549-77.

Saloner, G. and A. Shepard, 1995. Adoption of Technologies with Network Effects: An Empirical Examination of the Adoption of Automated Teller Machines. *RAND Journal of Economics*, 26, 479-501.

Wei, Chao, 2003. Energy, the Stock Market, and the Putty-Clay Investment Model. *American Economic Review*, 93, 311-24.

Wing, Ian Sue and Richard Eckaus, 2004. Explaining Long Run Changes in the Energy Intensity of the U.S. Economy. mimeo.

Table 1

Summary Statistics from the Census of Manufactures, 1967-1997

	1967	1972	1977	1982	1987	1992	1997
				Panel A: Entrants			
Log Energy Efficiency	-4.477	-4.458	-4.678	-5.329	-5.028	-4.889	-5.115
	(1.037)	(1.044)	(1.090)	(1.269)	(1.188)	(1.345)	(1.398)
Weighted Log Energy Efficiency	-5.009	-4.860	-5.032	-5.623	-5.329	-5.427	-5.578
	(1.156)	(1.135)	(1.287)	(1.580)	(1.609)	(1.682)	(2.032)
Energy Cost Share	0.026	0.022	0.022	0.026	0.024	0.026	0.025
	(0.583)	(0.504)	(0.124)	(0.290)	(0.448)	(0.313)	(0.893)
				Panel B: Incumbents			
Log Energy Efficiency	-4.612	-4.489	-4.687	-5.157	-4.963	-4.892	-5.122
	(0.997)	(1.029)	(1.093)	(1.255)	(1.176)	(1.292)	(1.380)
Weighted Log Energy Efficiency	-4.883	-4.688	-4.744	-5.060	-5.106	-5.134	-5.419
	(1.057)	(1.091)	(1.253)	(1.352)	(1.394)	(1.572)	(1.852)
Energy Cost Share	0.018	0.018	0.022	0.039	0.021	0.024	0.019
	(0.036)	(0.048)	(0.037)	(4.521)	(0.110)	(0.536)	(0.062)
				Panel C: Aggregate Statistics			
Entry Rate	0.23	0.33	0.29	0.28	0.25	0.29	0.26
Number of Plants	150,456	160,425	179,311	199,721	190,598	209,505	194,277

Variables are constructed from the Census of Manufactures (CM), 1967-1997. Each cell in Panels A and B report the mean, with standard deviation in parentheses, for plants in the corresponding sample. Panel A includes plants in the indicated Census year that did not appear in the previous Census. Panel B includes plants that operated in the previous Census. Energy use for each plant is computed as reported energy expenditure divided by the price of energy for the plant's state and industry. Output is the reported total value of shipments divided by the plant's output price deflator. All prices are normalized to one in 1967 (see text). Log energy efficiency is the log of the ratio of energy use to output. Weighted log energy efficiency uses plant shipments as the weight. Energy cost share is the ratio of energy expenditure to shipments. Entry rate is the number of entrants in the subsample, divided by the total number of plants.

Table 2

Response of Energy Efficiency to the Price of Energy, 1967-1997

	Baseline: industry x region x year, cohort and state controls	Baseline and industry x state	Baseline and industry x price of energy	Baseline and cohort x industry	Omit outliers
	Dependent Variable: Log Energy Efficiency				
	(1)	(2)	(3)	(4)	(5)
Log Price of Energy x Entrant	-0.099 (0.026)	-0.077 (0.044)	-0.095 (0.032)	-0.091 (0.030)	-0.082 (0.026)
Log Price of Energy	-0.172 (0.019)	-0.198 (0.023)	-0.248 (0.055)	-0.181 (0.047)	-0.200 (0.016)
Entrant	-0.083 (0.010)	-0.023 (0.012)	-0.082 (0.013)	-0.126 (0.037)	-0.027 (0.009)
R^2	0.72	0.74	0.74	0.74	0.79
Number of Observations	1,284,293	1,284,293	1,284,293	1,284,293	1,100,187

Huber-White standard errors in parentheses. The dependent variable is log energy efficiency, constructed as in Table 1. Log price of energy is the log of the weighted sum of state by year energy source prices from the DOE, using BTU weights constructed from the LRD, divided by the output price (see equation (13)). Entrant is a dummy variable, equal to one if the plant appears for the first time in the CM. Log price of energy x entrant is the interaction of the two variables. The sample is constructed as described in the text. All regressions are estimated by Weighted Least Squares, using the plant's share in shipments for the corresponding year as the weight. The cohort of a plant refers to the year the plant first appears in the Census. All regressions include a full set of cohort dummy variables, equal to one if the plant belongs to the corresponding cohort, and equal to zero the year the plant enters. All columns include a full set of industry-region-year interactions and state dummies, and column 2 also contains industry-state interactions. Column 3 includes interactions of the price of energy with the industry's 3-digit industry. Column 4 includes the interactions of the plant's entry cohort with its 2-digit industry. Energy cost share is constructed as in Table 1 and the 5th and 95th percentiles of the cost share are computed for each industry-region-year-entrant cell. Column 5 omits outlying observations.

Table 3

Effect of the Price of Energy on Entry, Survival of Entrants and Exit

	(1)	(2)	(3)	(4)	(5)	(6)
	Log entry	Entry fraction	Log survival of entrants	Survival probability of entrants	Log exit	Exit probability
Log Price of Energy	0.059 (0.027)	-0.003 (0.004)	0.057 (0.043)	0.006 (0.014)	0.118 (0.036)	0.009 (0.005)
R^2	0.89	0.62	0.88	0.57	0.89	0.63
Number of Observations	62,069	95,022	44,938	53,946	53,436	83,980

Huber-White standard errors in parentheses. Observations are at the industry-state-year level. The dependent variable is the log of a count of entering plants in column 1; the count of entering plants divided by the number of operating plants in the corresponding cell in column 2; the log of the number of entrants that survive to the next Census in column 3; the number of survivors divided by the number of entrants in column 4; the log of the number of plants in the previous Census that exit before the current Census in column 5; and the fraction of plants that exit between the previous and current Census in column 6. All regressions are estimated by Weighted Least Squares, with the share in total shipments for the corresponding year as the weight. All regressions include industry-region-year interactions and state dummies. Columns 1 and 2 include observations from 1967-1997, and columns 3-6 include 1967-1992.

Table 4

Controls for Plant Productivity

	Log entry	Entry fraction	Log survival of entrants	Survival probability of entrants	Log initial capital stock	Log initial employment
			Dependent Variable: Log Energy Efficiency			
	(1)	(2)	(3)	(4)	(5)	(6)
Log Price of Energy x Entrant	-0.093 (0.026)	-0.099 (0.026)	-0.111 (0.024)	-0.107 (0.024)	-0.117 (0.024)	-0.096 (0.027)
Log Price of Energy	-0.177 (0.020)	-0.173 (0.019)	-0.213 (0.022)	-0.211 (0.022)	-0.184 (0.019)	-0.168 (0.019)
Entrant	-0.085 (0.010)	-0.082 (0.010)	-0.077 (0.010)	-0.080 (0.009)	-0.051 (0.011)	-0.055 (0.010)
Plant Productivity Variable	-0.009 (0.005)	-0.012 (0.023)	-0.020 (0.005)	-0.045 (0.020)	0.045 (0.003)	0.046 (0.003)
R^2	0.71	0.72	0.68	0.69	0.74	0.73
Number of Observations	1,223,646	1,282,864	1,000,627	1,034,807	1,069,753	1,269,962

Huber-White standard errors in parentheses. The dependent variable, log price of energy and entrant variables are computed as in Table 2. The sample is constructed as in Table 2, except that column 5 does not include observations from 1967 or 1997. Regressions are estimated by Weighted Least Squares, using the share of plant shipments in the corresponding year as the weight. All regressions include the same control variables as in column 1 of Table 2. Plant Productivity Variable refers to the variable added to the regression that is indicated in the column heading. In columns 1-4 plant productivity is the dependent variable from columns 1-4 of Table 3, for the plant's industry, in the year and state it entered. Column 5 includes the log book value of the capital stock the year the plant entered (or the 1972 value for plants that entered before 1972). Column 6 uses the log of the plant's employment in the first observed year of operation.

Table 5

Substitution Elasticities of Entrants and Incumbents

	Log price of energy x new plant interactions	Baseline specification, with log price of energy x year interactions	Log price of energy x region x year interactions	Log price of energy x industry x year interactions
		Dependent Variable: Log Energy Efficiency		
	(1)	(2)	(3)	(4)
Log Price of Energy x Entrant		-0.098 (0.026)	-0.099 (0.026)	-0.096 (0.026)
Log Price of Energy	-0.294 (0.040)	-0.308 (0.118)	-0.292 (0.119)	-0.299 (0.118)
Entrant		-0.078 (0.012)	-0.073 (0.032)	-0.071 (0.032)
Log Price of Energy x 1972 New Plant	-0.018 (0.016)			
Log Price of Energy x 1977 New Plant	-0.007 (0.031)			
Log Price of Energy x 1982 New Plant	0.040 (0.042)			
Log Price of Energy x 1987 New Plant	-0.070 (0.044)			
Log Price of Energy x 1992 New Plant	0.043 (0.049)			
Log Price of Energy x 1997 New Plant	0.021 (0.067)			
R^2	0.73	0.72	0.73	0.73
Number of Observations	1,133,837	1,133,837	1,133,837	1,133,837

Huber-White standard errors in parentheses. The dependent variable, log price of energy and entrant dummy are computed as in Table 2. The sample is the same as column 1 of Table 2, except that it does not include observations in 1967. Regressions are estimated by Weighted Least Squares, using the share of plant shipments in the corresponding year as the weight, and include industry-region-year interactions and state dummies. Columns 2-4 include cohort dummies. Operate is a set of dummy variables, equal to one if the plant operated in the corresponding year. New plant is a set of dummy variables, equal to one if the plant entered that year. Column 1 includes the operate and new plant variables, plus the interaction of all variables with the log price of energy. Column 2 includes a full set of log price of energy by year interactions. Column 3 includes log price of energy by region by year interactions, and column 4 includes log price of energy by 2 digit industry by year interactions.

Table 6

Alternative Measures of the Price of Energy

	Log aggregate price	Log aggregate forecasted price	5-year lag	5-year average	NBER prices
			Dependent Variable: Log Energy Efficiency		
	(1)	(2)	(3)	(4)	(5)
Log Price of Energy x Entrant	-0.045 (0.012)	-0.122 (0.020)	-0.098 (0.026)	-0.102 (0.027)	-0.061 (0.035)
Log Price of Energy			-0.162 (0.019)	-0.162 (0.019)	
Entrant	-0.040 (0.012)	-0.018 (0.012)	-0.087 (0.010)	-0.088 (0.010)	-0.068 (0.009)
R^2	0.72	0.72	0.72	0.73	0.70
Number of Observations	1,284,293	1,284,293	1,277,911	1,241,033	1,090,016

Huber-White standard errors in parentheses. The dependent variable and entrant dummy are computed as in Table 2. The sample is constructed as in Table 2, except that column 5 does not include observations in 1997. Regressions are weighted by plant shipments, estimated by Weighted Least Squares, and include the same control variables as in column 1 of Table 2. Column 1 uses the log aggregate price of energy, the log of the weighted mean of the individual plant prices. The log aggregate forecasted price is calculated from a VAR with two lags, estimated using the DOE state prices for natural gas, electricity, coal, residual and distillate. The forecasted price of each energy source is the sum of the current price and the linear prediction of the price over the next ten years, discounted by ten percent. The log aggregate forecasted price of energy is the log of the weighted sum of the energy source forecasts, using aggregate weights (see text). Column 2 includes the interaction of the log forecasted price and the entrant dummy. Column 3 uses the 5-year lag of the log price of energy for the plant's industry and state. Column 4 uses the mean log price of energy for the plant's industry and state over the previous 5 years. Column 5 uses the industry by year log prices of energy and shipments from the NBER Manufacturing Productivity Database.

Table 7

Additional Robustness Checks

	Cluster standard errors by year	Cluster standard errors by plant	Dep var is log fuel intensity	Dep var is log electricity intensity	Include dummy if plant is part of a new firm	Omit top quartile of energy intensive industries	Omit electricity producing industries	Separate effect for coal intensive industries	Omit low natural gas counties	Dependent variable is value added
					Dependent Variable: Log Efficiency					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Log Price of Energy x Entrant	-0.099	-0.099	-0.112	-0.116	-0.073	-0.115	-0.112	-0.096	-0.101	-0.100
	(0.016)	(0.024)	(0.034)	(0.024)	(0.032)	(0.031)	(0.029)	(0.033)	(0.027)	(0.030)
Log Price of Energy	-0.172	-0.172	-0.175	-0.224	-0.173	-0.168	-0.158	-0.173	-0.173	-0.183
	(0.051)	(0.023)	(0.021)	(0.022)	(0.019)	(0.022)	(0.020)	(0.022)	(0.019)	(0.021)
Entrant	-0.083	-0.083	-0.085	-0.025	-0.117	-0.075	-0.088	-0.093	-0.083	-0.015
	(0.019)	(0.011)	(0.013)	(0.011)	(0.015)	(0.012)	(0.011)	(0.010)	(0.011)	(0.012)
R^2	0.72	0.72	0.75	0.65	0.72	0.66	0.68	0.72	0.72	0.71
Number of Observations	1,284,293	1,284,293	913,574	1,175,548	1,284,293	980,526	1,152,999	1,171,460	1,164,375	1,276,976

Huber-White standard errors in parentheses. The dependent variable in columns 1, 2 and 5-10 is log energy efficiency, constructed as in Table 2. The dependent variable is log fuel efficiency in column 3 and log electricity efficiency in column 4, constructed using reported fuel and electricity expenditure and the real price of fuels and electricity (see text). Specifications in columns 1, 2 and 5-10 include the same independent variables as in column 1 of Table 2, and observations are weighted by plant output. Column 3 uses the real price of fuel in place of the real price of energy to construct the independent variables, and column 4 uses the real price of electricity (see text). In column 1 standard errors are clustered by year, and by plant in column 2. In column 5 a dummy variable, equal to one if the plant does not belong to a pre-existing firm, is added plus its interaction with the price of energy. Column 6 omits plants belonging to industries in the top quartile of energy cost share, computed from the 1963 Census. Column 7 excludes plants in the four 2-digit industries with the largest share of electricity sales in total output: Paper, Chemicals, Petroleum and Primary Metals (see text). Coal cost share is the ratio of coal expenditure to total energy costs by industry. Column 8 includes the interactions with the main independent variables of a dummy equal to one if the industry is above the 90th percentile of coal cost share. Column 9 excludes counties with natural gas consumption below the 15th percentile (see text). Real value added is equal to gross shipments, net of materials and energy expenditure, divided by the value added deflator. In column 10 the dependent variable is the log of real energy use divided by real value added, where real energy use is computed as in Table 2.

Table 8

Investment and Capital Retirements of Existing Plants, Using a Balanced Panel

|	(1)	(2)	(3)	(4)	(5)	(6)
	Log investment	Log investment with 5-year lag price of energy	Log investment with forecasted price of energy	Log retirements	Log retirements with 5-year lag price of energy	Log retirements with forecasted price of energy
Log Price of Energy	-0.020 (0.017)	-0.038 (0.018)		0.324 (0.158)	0.257 (0.177)	
Log Price of Energy x Energy Intensity	-0.055 (0.522)	0.990 (0.477)	-0.204 (0.371)	-7.572 (4.131)	-3.134 (4.146)	-11.913 (6.659)
R^2	0.24	0.23	0.24	0.27	0.27	0.27
Number of Observations	180,210	178,785	180,210	141,933	140,855	141,944

Huber-White standard errors in parentheses. The sample includes a balanced panel of plants appearing in all CM and ASM years from 1972-1988. All variables are in first differences. All regressions include industry-region-year dummies and are estimated by Weighted Least Squares, using the share of plant shipments in the corresponding year as the weight. The dependent variables are indicated in the column headings. Plant capital stocks are constructed using the perpetual inventory method, beginning with the 1972 book value of capital, subtracting retirements and deflating investment (see text). Log investment is the difference between the log capital stock in the current and previous years. The dependent variable in columns 1-3 is log investment, and in columns 4-6 it is log retirements. Columns 1 and 4 use the log current price of energy, columns 2 and 5 use the 5-year lag, and columns 3 and 6 use the log aggregate forecasted price of energy, using a VAR with two lags (see Table 6 and the text). Energy intensity is the ratio of total energy costs to output in 1963 for the plant's industry, calculated from the 1963 Census of Manufactures.

## Figure 1: Substitution Elasticity vs Technological Change
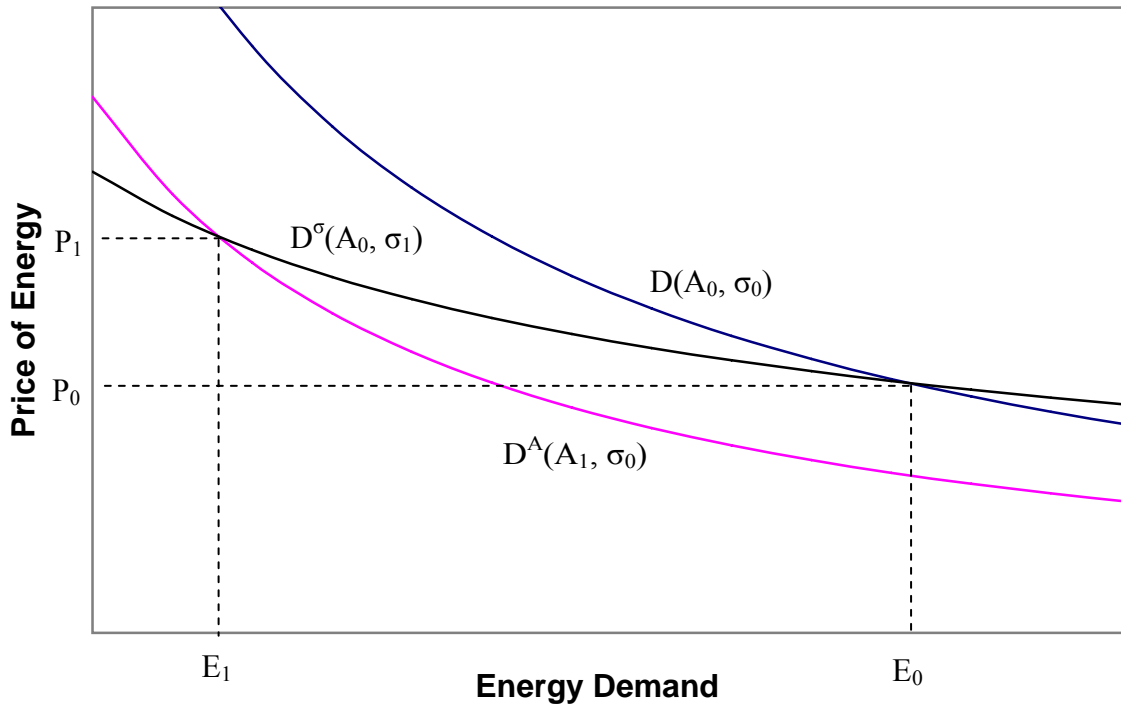
The curves represent the demand for energy, as functions of the energy technology, A, and the price elasticity, σ. The price of energy is P_0 before a price shock, and after the shock the price is P_1. The initial energy consumption is E_0 and final consumption is E_1. The curve $D^{\sigma}(A_0, \sigma_1)$ has the same energy technology and a greater price elasticity than $D(A_0, \sigma_0)$. The curve $D^A(A_1, \sigma_0)$ has the same price elasticity and a different energy technology from $D(A_0, \sigma_0)$.

Figure 2: Real Energy Prices, 1970-1997
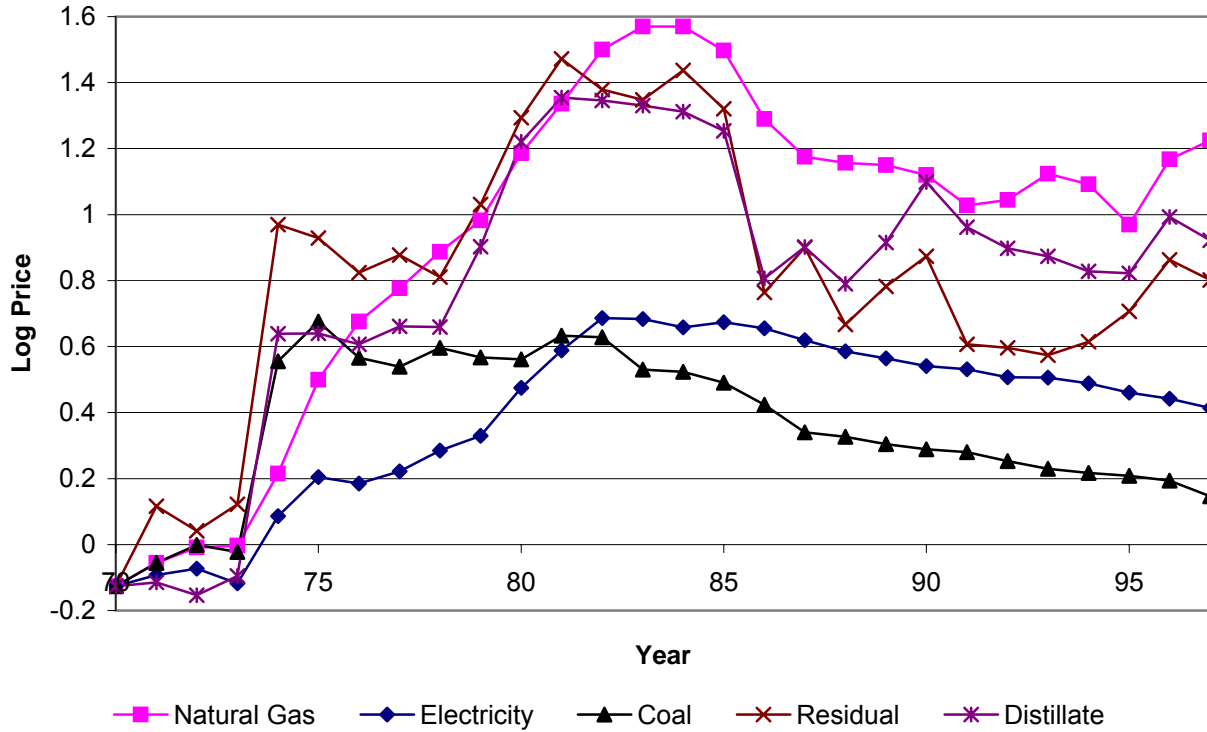
The real price of each energy source is the nominal price, from the State Energy Price Report, 1970-1997, divided by the output price, from the Census of Manufactures, 1967-1997 (see text). All prices are normalized to one in 1967.

Figure 3: Mean Energy Efficiency of Entrants and Incumbents, 1972-1987

Incumbents include plants that operated in the previous Census. Each entrant cohort includes plants that entered in the indicated year, and are still operating in the corresponding year. Each data point is the weighted mean log energy efficiency, computed as in Table 1, using plant shipments as the weight (see text).