Is There An Energy-Efficiency Gap? Experimental Evidence from Indian Manufacturing Plants*

Nicholas Ryan†

March 31, 2015

Abstract

Informational market failures have been faulted as one reason why unproductive firms do not adopt profitable technologies, and energy-saving technologies in particular. This paper studies whether such informational market failures reduce the energy-efficiency of Indian manufacturing plants using a large randomized-controlled trial that offers information, via industrial energy audits, and skilled labor, via continuing energy consultancy, to encourage energy-efficiency investments. Energy audits project the returns to energy-efficiency to be very high, forecasting total savings of ten percent of plant energy bills on a collection of small investments. I find that, despite these projections, energy audit treatment plants invest marginally and insignificantly more than control plants and far less than was forecast to be profitable. This investment appears to slightly increase plant physical efficiency, as measured in independent technical surveys, and treatment plants increase capacity utilization in response. Consistent with this increase in use, electricity demand temporarily and insignificantly declines in the treatment, then rebounds to the control level. The overall pattern of low adoption is consistent with ancillary fixed costs of adoption that make small-scale but seemingly high return investments unprofitable.

JEL Codes: O14, Q41, D24, L65, L67

1 Introduction

Are firms in developing countries unproductive due to market forces, or market failures? There are huge differences between the productivities of developing-country firms competing in the same industries, and more productive firms, which should displace their rivals, instead grow at a sluggish pace (Hsieh and Klenow, 2009, 2014). These differences may arise because firms are making optimal investment decisions in an economic environment that is more volatile or less competitive than that in developed countries.[1] Alternatively, productivity differences

*I thank Esther Duflo, Michael Greenstone and Rohini Pande for continual support and guidance. I also thank R N Pandya of the Gujarat Energy Development Agency for advice and encouragement. Harsh Singh, Harsh Vijay Singh, Vipin Awatramani, Raunak Kalra and Maulik Chauhan provided exemplary research assistance and project management. I thank the MIT Energy Initiative, Veolia Environment, US AID – Development Innovation Ventures and the Sustainability Science Program at Harvard for financial support. All views and errors are my own.

†Yale University, Dept. of Economics, Box 208269, New Haven, CT 06520-8269, nicholas.ryan@yale.edu

[1] Asker, Collard-Wexler and Loecker (2014) argue that nearly all of the cross-country dispersion in productivity can be explained by a dynamic model where firms face capital adjustment costs and more variable

may be due to market failures that distort the allocation of inputs across firms. For example, Bloom et al. (2013) argue that firms do not adopt profitable changes in management practices because of informational barriers or market failures.

Exactly such market failures have long been blamed for low levels of investment in energy-efficiency, with the phrase "energy-efficiency gap" coined to denote differences between the privately optimal and actually observed levels of energy-saving investment (Jaffe and Stavins, 1994; Allcott and Greenstone, 2013). The claim is not that energy consumption causes externalties and therefore private investment in saving energy is less than socially optimal. Rather, it is stronger: that investment is less than *privately* optimal, based only on future savings in private energy costs, because market failures inhibit firms or consumers from investing enough. The most important potential market failures regard information, especially that agents are poorly informed and unable to measure or value future energy savings. This lack of information may persist, even when suppliers of efficient products have an incentive to educate energy consumers, due to difficulties in contracting on realized energy savings.[2]

This paper studies whether such informational market failures reduce the energy-efficiency of Indian manufacturing plants using a large randomized-controlled trial that offers information and skilled labor to encourage energy-efficiency investments. The sample of over four hundred plants was recruited from the chemical and textile sectors in industrial clusters of the Indian state of Gujarat. The stakes for these plants to save energy are high, as the average plant spends USD 200,000 on electricity and fuel each year, about 15-20% of total costs. The study worked with Department of Climate Change, Govt. of Gujarat, to identify leading energy consultants in the state and enlist them to offer services to a random subset of these plants. The services, offered as experimental treatments, were of two kinds. First, a random half of plants were offered information in the form of detailed energy audits, a several-day review of specific energy consumption in the plant that looks for profitable investments to reduce energy bills. The output of this review is an energy audit report, presented to the owner or plant manager, that proposes investment options and details the required capital

productivity shocks. In the United States, (Syverson, 2004) argued that transport costs can explain the degree of competition and hence productivity dispersion in an oligopoly market; it is plausible that, with higher transport costs (Atkin and Donaldson, 2014), this effect would be greater in developing countries.

[2]There is a large market in the United States for Energy Service Companies or ESCOs, which are private energy consultants that advise firms on saving energy, put up funds to make capital improvements, and then are paid with a share of the reduction in energy bills after the fact. This type of contract may lead to hold-up if the client cannot be forced to share the energy savings, which problem may be worsened when counterfactual energy consumption is difficult to estimate.

and projected energy bill savings for each option. Energy audits offered in this experiment were subsidized by the state Government under a program that predated the study. Second, a random half of those plants completing energy audits were offered skill in the form of energy managers, engineers who would visit the plant periodically after the audit to encourage adoption of its recommendations, supervise procurement and installation of more efficient equipment and advise plant staff on operating practices.

The data used to study plant efficiency and the results of these interventions are severalfold. Prior to the assignment of treatment a brief baseline survey was conducted to collect plant characteristics such as energy consumption and sales and decide eligibility to enter the sample. In the course of the treatments, consultants provided energy audit reports to the research team that list all recommendations they made for energy-saving measures, along with how much each measure costs up-front and how much it is expected to save in energy costs. Then, at least twelve months after the completion of energy audits, the research team and expert engineers—wholly independent of the consultants offering their services in the treatment—returned to plants to record production details, investments in equipment upgrades and maintenance, and physical measurements of the efficiency of energy-using systems within the plant. The survey collects energy consumption data and additionally gets plant consents to obtain their electricity bills directly from their electric utility. Production theory, described in Section 2 below, shows that it is critical to measure these components, like inputs, efficiency, output and prices, separately, since the plant may respond to improvements in energy-efficiency by substituting towards greater energy use in production, offsetting some of the projected energy savings.

The paper has four sets of findings with regard to projected savings in energy audits and then actual investment, physical efficiency and energy and other input consumption in response to the treatments. First, consultants project high returns to energy-efficiency investment, with a median annual return on recommended measures of 104%. The investments required for these measures are generally small, with a median of USD 361; 95% of individual measures require investment of less than 1.7 percent of plant capital stock.[3] The total scope of investment is also limited, with the investment required to adopt all recommendations totaling less than three percent of baseline capital stock, however, because returns are projected to be

[3] All monetary values in the paper were originally in Indian rupees and have been coverted to dollars at an exchange rate of USD 1 = INR 45.

so high, this degree of investment is forecast to save 11% of plant energy bills.

Second, despite these high returns, plants assigned to the energy audit treatment are estimated to invest a modest and statistically insignificant USD 368 (standard error 256, p-value 0.15 against the null of zero additional investment) more in energy-saving equipment upgrades and maintenance. To get a sense of how this compares to projections, suppose that plants had a marginal return to capital of 105%, as found for mid-size Indian firms in Banerjee and Duflo (2014), and adopted energy audit recommendations that were projected to have a return equal or greater than this rate. This decision rule would yield USD 6,981 of investment per energy audit treatment plant; we can therefore easily reject that firms take projections at face value and invest with this hurdle rate (p-value < 0.01), or significantly higher rates in mind.

Third, in the endline survey's direct meaures of efficiency, energy audit treatment plants are found to be insignificantly more efficient overall but measurably more efficient in some plant systems. The survey encompassed a large number of efficiency measures from different areas of the plant. We aggregate these into a single efficiency index and find that treatment plants are 0.04 standard deviations (standard error 0.06 standard deviations) more efficient than control plants, but, for the boiler system, an important determinant of energy consumption, are a significant 0.29 standard deviations (standard error 0.17) more efficient. Plant capacity utilization overall is estimated to increase by a statistically significant one hour per day on a base of 17 hours per day in the control, and this increase is largest for the boiler system, which saw the greatest increase in efficiency. Most of these plants are running at less than capacity; this evidence suggests that modest increases in efficiency due to energy audits led to increases in capacity utilization.

Fourth, energy consumption does not significantly decline in the energy audit treatment group relative to the control group. The best measure of energy consumption is electricity consumption, because electricity is metered by and data obtained from electric utilities. Electricity consumption appears to decline in the months after an energy audit, but I estimate that on average over the first year electricity consumption declines an insigificant 2,103 kWh per month (standard error 2,436 kWh per month), or 3.7% on a control mean of nearly 57,000 kWh per month. In the second year after energy audit, electricity demand is estimated to actually increase in the treatment group, relative to the control, by a comparable amount.

These estimates are fairly precise, but due to the low level of additional investment it is not possible to reject either the hypothesis that realized energy savings in the first year are zero or that realized energy savings are equal to projections for the amount of investment undertaken. The pattern of lower savings over time is consistent both with the theoretical prediction that energy savings should decline in the longer run, as other inputs and output can adjust to energy efficiency, and with the empirical finding of higher plant hours of operation.

Why do plants not invest more, given such high projected savings? I consider several hypotheses to explain low adoption and conclude that the data is consistent with ancillary fixed costs of adoption that make high-return but small investments unprofitable. Notably, while the average energy audit treatment effect is an insignificantly positive USD 367, the 95th-quantile treatment effect is estimated to be far larger USD 72 thousand (p-value 0.07 against the hypothesis of no change in investment). This large effect in the upper tail is consistent with plants having infrequent opportunities to significantly adjust their capital stock, and responding, in those opportunities, to the recommendations offered in energy audits.

This paper contributes to the literatures in energy economics on the adoption of energy-efficient technology and in development economics on the returns to entrepreneurship and productivity dispersion. In energy economics, the idea of an energy-efficiency gap began with studies showing that moderate or high discount rates were needed to justify observed choices of efficiency (Hausman, 1979; Train, 1985). This idea has been taken up to support an active policy stance of energy-efficiency subsidies and mandates, but credible evidence that market failures contribute to any gap remains thin to non-existent (Metcalf and Hassett, 1999; Allcott and Greenstone, 2012). Recent research has begun to credibly measure the returns to energy-efficiency investment and has generally found that returns to efficiency investments may be overstated.[4] This paper is the one of a very few experimental studies in the literature and, to my knowledge, the only causal study of energy-efficiency investment by firms.[5] The findings

[4](Davis, Fuchs and Gertler, 2014) find that energy savings from more efficient residential appliances are far less than projected, due in part to more intensive consumer use of more efficient appliances. Fowlie, Greenstone and Wolfram (2015) find low demand for even free residential weatherization investments in Michigan, and, in ongoing work, that this may be due to lower-than-projected energy savings. Jacobsen and Kotchen (2010), to the contrary, finds that projected and realized energy savings for building code modifications in Florida line up well.

[5]Anderson and Newell (2004) study industrial energy audits by the U.S. Department of Energy's Industrial Assessment Centers and find take-up of 53% of audit recommendations, but are not able to observe energy consumption or physical efficiency after audits, or to establish a counterfactual for adoption. Rohdin, Thollander and Solding (2007) report that Swedish foundris have greater trust in information provided by their association than in government-sponsored energy audits.

here agree with several recent studies of households in finding that widespread market failures are not needed to explain low investment in energy-efficiency.

In development economics this paper also contributes to a literature on using firm-level experiments, often via consulting or other informational interventions, to understand the constraints on micro-enterprise growth. This literature has shown mixed results for the effect of training or consulting on firm profits for micro-enterprises, with some interventions increasing, and others decreasing, average profits (Fischer and Karlan, 2015; Drexler, Fischer and Schoar, 2014; Karlan, Knight and Udry, 2012). This paper provides a rare study on larger enterprises with significant capital investment and employment, and suggests that investment decisions in these enterprises are relatively insensitive to consulting.

The closest antecedent to this study is Bloom et al. (2013), which is an experimental study of management consulting in a sample of large textile mills in India. Bloom et al. (2013) find large effects of this consulting on management practices, quality and productivity, and estimate that these effects increase profits by more than the cost of consulting, even in the first year of adoption. The present study finds small effects of energy consulting and weak take-up of energy-saving measures projected to have high returns, suggesting that informational market failures are not the main constraint to adoption in this domain. Fixed costs in technology adoption can parsimously explain the findings of the current paper, as well as the differences between these experiments, since the design of the Bloom et al. (2013) experiment was to offer much more intensive consulting for a far smaller number of plants.

The rest of the paper runs as follows. Section 2 models how plant energy use may respond to energy-efficiency improvements depending on the flexibility of other factors of production. Section 3 gives background on the setting for the study and the design of the experiment. Section 4 describes the empirical results, and Section 6 concludes by discussing the relevance of the findings for energy-efficiency policy.

2 Model of Plant Response to Energy-Efficiency

Consider a plant that produces physical output (e.g., meters of textile) with the constant elasticity of substitition production function $Q = [(A_X X)^\rho + (A_E E)^\rho]^{1/\rho}$, where X is a composite factor of production, including factors such as capital and labor, and E is energy consumption. Each factor has a factor-specific efficiency, so that A_E is energy-efficiency and $A_E E$ the input of energy services. The elasticity of substitution between E and X is $\sigma = \frac{1}{1-\rho}$.

Let the price of output be p, the price of energy p_E and the price of the composite factor be one. Assume that the plant faces a downward-sloping demand for its output $Q = Bp^{-\epsilon}$.[6] Then the plant's revenue-production function is

$$Y(X, E) = B^{1/\epsilon} \left[(A_X X)^\rho + (A_E E)^\rho) \right]^{\phi/\rho}.$$

Where $\phi = \frac{\epsilon - 1}{\epsilon}, \epsilon > 0 \Rightarrow \phi < 1$, so that the production of revenue has decreasing returns to scale, and absolute input demands are well defined. Define the per-period profit function given efficiency A_E as

$$\Pi(A_E) = \max_{X,E} B^{1/\epsilon} \left[(A_X X)^\rho + (A_E E)^\rho) \right]^{\phi/\rho} - X - p_E E$$

What is meant by an energy-efficiency gap? Suppose that the plant first chooses efficiency and then produces with that efficiency thereafter, facing static demand in each period. In starting production plants can invest in efficiency using capital and skill, $A_E = f(K_E, H_E)$, where $f(\cdot, \cdot)$ is continuous, increasing and concave in both its arguments. The optimal choice of efficiency is governed by $A_E = f(K_E^*, H_E^*)$, where

$$(K_E^*, H_E^*) = \arg\max_{K_E, H_E} \frac{\Pi(A_E)}{1-\delta} - K_E p_K - H_E p_H.$$

We define an energy-efficiency gap with respect to capital and skill as

$$\frac{\frac{\partial \Pi(A_E)}{\partial A_E} \frac{\partial f}{\partial K_E}}{1-\delta} > p_K \qquad \frac{\frac{\partial \Pi(A_E)}{\partial A_E} \frac{\partial f}{\partial H_E}}{1-\delta} > p_H. \tag{1}$$

These conditions state that the present discounted value of the product of the marginal profits from efficiency and the marginal efficiency from investment in a factor that produces efficiency exceed the price of that factor. That is, a gap is defined as when improving the efficiency of one's plant increases profits more than it costs today. Whether a gap exists will therefore depend on the appropriate discount factor δ as well as the effect of efficiency on profits.

[6] This assumption is common, often with relatively high ϵ to represent competitive industries (Foster, Haltiwanger and Syverson, 2008; Allcott, Collard-Wexler and O'Connell, 2014).

A simpler criteria to define a gap would be based on energy savings alone. Suppose that $p_E \frac{\partial f / \partial K_E}{1-\delta} > p_K$, meaning that the present value of energy savings is greater than the price of investment in capital needed to achieve those savings. This formulation is sufficient but not necessary for a gap, as defined in (1) above, since it assumes that no other factors of production adjust to energy-efficiency. If other inputs or output can adjust, then energy savings may be small or negative, even when a gap exists and investment in efficiency is worthwhile.

Consider the response of energy consumption to changes in energy-efficiency. First, suppose that energy services are fixed, $A_E E = \overline{S}$. Then $\log E = \log \overline{S} - \log A_E$ so that the elasticity of energy consumption with respect to efficiency is $\varepsilon_{E,A_E} = -1$. This case of no input adjustment is what energy audit projections assume; the change in profits with efficiency is equal to the price of energy inputs, since energy consumption moves inversely with efficiency. This is a reasonable case for small improvements of efficiency and in the short-run, where we may expect energy inputs to be flexible from day to day–such as via drawal from the electricity grid–but labor contracts, materials orders and certainly capital to adjust more slowly.

In the medium-run, suppose the quantity of production $Q = \overline{Q}$ is fixed but demand for other inputs is flexible. Then we may derive energy demand and calculate the elasticity of demand with respect to efficiency as

$$\varepsilon_{E,A_E}|_{Q=\overline{Q}} = -1 + \sigma \frac{(A_X p_E)^{\sigma-1} e^{1-\sigma} A_E}{(A_X p_E)^{\sigma-1} A_E^{1-\sigma} + 1}.$$

Note $\sigma \geq 0$ and the right-hand term is positive, so that energy use is found to decrease less than in the short-run case, as the plant substitutes towards E consumption depending on the value of σ. In the case of Leontief production $\sigma = 0$ and efficiency again reduces energy consumption one-for-one; for any positive degree of substitution the energy savings will be muted.

Now return to a longer-run case where quantity $Q = Bp^{-\epsilon}$ can adjust. I derive the elasticity of energy demand with respect to efficiency as

$$\varepsilon_{E,A_E} = \varepsilon_{E,A_E}|_{Q=\overline{Q}} + \epsilon - \left(\sigma \frac{\epsilon - 1}{\epsilon} + 1 \right) \frac{(A_X p_E)^{\sigma-1} e^{1-\sigma} A_E}{(A_X p_E)^{\sigma-1} A_E^{1-\sigma} + 1}.$$

Thus, in addition to the substitution effect present in the fixed-quantity case, energy consumption can additionally increase due to an expansion governed by the elasticity of demand $\epsilon > 0$. In some cases this could be large, such as for large ϵ but small σ, when the more-efficient plant can expand by offering a lower price but does so mainly with energy use, since inputs are not very substitutable.

There are several takeaways from this theory on plants' response to improved efficiency that will guide the empirical work. First, the correct measure of energy-efficiency gap is a comparison of change in profit to investment, not only the discounted value of energy consumption. Therefore the empirical analysis will encompass energy consumption but also pricing and other input responses to the experimental treatments.[7] Second, in the short run prior to the response of other inputs, energy consumption is sufficient to measure the plant response to improvements in efficiency. The energy consumption response is expected to be largest in the short run since the input of energy services is roughly fixed. Third, in the medium and longer runs greater efficiency will induce substitution towards energy service as an input and some expansion of production, which will offset the initial reduction in energy consumption to a degree that depends on the elasticities of substitution and demand.

3 Context and Experimental Design

(a) Energy-efficiency policy in India

Energy is a productive input of special importance from a policy view. From 2010 to 2040, energy use in non-OECD countries is projected to increase 90 percent, as compared to 17 percent in the OECD (Energy Information Administration, 2013). The industrial sector accounts for the largest share of energy demand and it is likely that 80% of energy use overall in 2040 will still come from fossil fuels. The combustion of these fuels causes large externalities through local air pollution that damages human health and greenhouse gas emissions that contribute to global climate change. Policy-makers have lately stressed the importance of efficiency, or a reduction in energy demand, for climate change mitigation. The head of the U.N. Climate Change Secretariat recently hailed energy efficiency as "the most promising

[7] Actual or imputed profits are difficult to measure since many plants decline to disclose production data and raw material consumption, which are targets of environmental regulation. Bloom et al. (2013) also report not being able to obtain accounting measures of profit despite deep engagement with sample firms.

means to reduce greenhouse gases in the short term."[8]

The belief underlying this statement is that energy-efficiency is cheap: if market failures make plants use energy inefficiently to begin with, fixing these failures can both reduce carbon emissions and save enough money on energy bills to be profitable even privately (See Allcott and Greenstone, 2012, for a review of the evidence on this idea). If this is true, a small amount of public investment in information or subsidies to technology adoption could yield large returns in lower energy consumption. Energy policy has been shaped around this idea. The United Nations Framework Convention on Climate Change (UNFCCC) has founded a Green Climate Fund to send money from developed to developing countries partly for climate change mitigation measure like energy-efficiency. In India, the US Agency for International Development, the Japanese International Cooperation Agency and the German overseas aid agency KfW are all active in industrial energy efficiency, having supported energy audits, technology development or subsidized lending for technology adoption.

The Indian government itself places a high priority on energy-efficiency in manufacturing. The government imposes a modest coal tax of INR 50 (approximately USD 1) per ton to fund clean energy technology, but the primary policy instruments for energy efficiency are informational and capital subsidies. The Bureau of Energy Efficiency, Ministry of Power has launched a "National Mission on Enhanced Energy Efficiency" across many sectors. For industry, this mission includes both an energy-conservation credit trading system for very large plants, and, for smaller plants, a nationwide campaign of energy audits and capital subsidies to identify energy-efficient technologies and encourage their wide adoption. This experiment was undertaken jointly with the Gujarat Energy Development Agency (GEDA), Department of Climate Change, Government of Gujarat, which is the BEE's state partner agency in the state of Gujarat.

(b) Selection of sample plants

The sample of plants was drawn from industrial associations with members in the textile and chemical sectors in the Indian state of Gujarat, which is home to 5% of India's population but 17% of industrial investment. In the chemical and textile processing sectors plants may spend 15-20% or more of their total production costs on energy. The Gujarat Energy Development

[8] De Boer, Yvo (August 28, 2007), available at http://www.reuters.com/article/idUSL2836333720070828 (last accessed August 8, 2013).

Agency (GEDA), a state government body responsible for the promotion of energy-efficient and renewable energy technologies in Gujarat, reports technically feasible “savings potentials” of around 20% of total energy bills for small plants in many energy-intensive sectors, including these two (GEDA, 2009). A recent Bureau of Energy Efficiency (BEE) study found that small chemical plants in Gujarat with less than 200 tons of production each year use 4050 kCal of energy per kilogram of product, 22% more than the 3312 kCal used by large plants (BEE, 2010).

A target sample size of 400 industrial plants was set to detect an 8% drop in electricity consumption with 80% statistical power, based upon energy consumption data from a sample of energy audits carried out by the Bureau of Energy Efficiency (BEE) for chemical factories in Ahmedabad, Gujarat. To reach this sample size, randomly selected industrial association members were assigned to be solicited, by energy consultants, for their interest to receive free energy consultancy, possibly including a detailed energy audit. A total of 925 plants were contacted, of which 53% said they were interested. Appendix Table A1 lists the reasons why firms declined, which typically did not relate to energy use *per se*: only 4% of plants said they already had an energy consultant, 4% that energy was not a large cost for their plant, and a further 5% that they expected the scope of savings was not large. Most plants that gave a reason for declining cited concerns about data confidentiality. From the 490 plants that responded with interest, the sample was cut down to 435 based on a maximum threshold for electricity load, in order to limit the sample to smaller plants and reduce the variance of energy demand in the sample.[9]

What kind of plants decided to enter the sample? In order to understand sample selection, I collected administrative data on industrial registrations from the Industries Commissioner, Government of Gujarat. Registration data includes details such as the capital stock and employment of plants, but has several limitations: it is only available for plants that register, is typically out of date, and it may be distorted if plants do not report truthfully to the government. Appendix Table A2 compares plant characteristics, in the registration data, for plants that were interested in the experiment versus not, among the 206 solicited plants that could be matched to this data set. The rate of interest is higher amongst these plants, at 75% instead of 53%, presumably because registered plants tend to be larger. Within the matched

[9] This restriction also has a policy motivation, in that most subsidized energy-audit programs restrict eligibility based on a maximum threshold for electricity load.

plants, most observable characteristics are similar, but interested plants tend to have a larger total capital stock, by USD 101 thousand (standard error USD 63 thousand). This difference is consistent with more capital-intensive plants selecting into the sample because they have more to gain from energy audits.

(c) Experimental treatments

The research design was a randomized-controlled trial with two intervention arms, meant to provide *information*, via plant energy audits, and *skill*, via energy managers. These treatments were chosen to test the leading hypotheses for why firms do not adopt energy-efficient technologies.

Energy audit treatment. A random half of treatment plants were offered free energy audits. An energy audit is a thorough, on-site review of how a plant uses energy and how it might profitably use less. Energy consultants employ electrical, chemical and mechanical engineers who spend approximately 6 man-days on site, depending on the size of the plant, collecting energy consumption information and measuring the efficiency of energy-using systems like motors, the boiler and the steam distribution system. At the conclusion of this measurement work, the consultant prepares an audit report suggesting investments to improve the efficiency of energy use, prioritized by their projected economic return. These reports are presented to the owner or plant manager of the audited plant in person, usually within two weeks of the completion of site work. As part of the experiment, the reports were also submitted to the research team and the Gujarat Energy Development Agency (GEDA), a co-sponsor, and both research and GEDA staff had the option to attend the presentation of reports.

Energy manager treatment. A random half of plants that completed energy audits and are interested in implementation are offered an energy manager to help in implementing audit recommendations. An energy manager is an engineer deputed to visit the plant for approximately 12 man-days over the course of several months, as decided jointly with the plant owner. This energy manager is responsible for identifying the most promising audit recommendations, procuring equipment, overseeing installation and training plant staff on any equipment or process changes.

The treatments were carried out by eight leading private energy consultants in the state of

Gujarat and the neighboring state of Maharashtra. The Gujarat Energy Development Agency (GEDA) had a program to subsidize industrial energy audits that pre-dated this study. Under this program, plants using a GEDA-certified consultant could get a subsidy of 50% of the cost of audits, up to a cap of INR 20,000 (about USD 450), which typically bound. The energy consultants participating in the study were solicited from those certified in 2009 and 2010 by GEDA. GEDA certifies 30 to 40 consultants as able to conduct thermal and electrical energy audits, which allows consultants to participate in the subsidy program as well as other government-sponsored consultancy and training activities. The consultants working in the study were deliberately selected from this group to be high-performing: the research team vetted consultants, in person and with the recommendations of GEDA, and invited eight of the best to conduct the project treatments on the basis of their reputations and past energy audit portfolio.

These services were paid for with a combination of research funds and government subsidies under GEDA's subsidy program. The total rate varied by consultant and plant between USD 900 to USD 1450. This total payment included USD 450 paid by GEDA on completion of the audit report and plant electricity bills. Research funds paid the rest of the total in equal installments on the completion of site work and the submission of the audit report to the plant. For the energy manager treatment consultants were paid at a flat rate of USD 800 to USD 1000 for the energy manager treatment, in two installments on the submission of progress and final reports.[10] The progress report was submitted after an initial meeting with the plant owner or manager to set out priorities and the final report to record any installations or upgrades that were done.

Table 1 summarizes the experimental design and the implementation of the treatments. Panel A shows the design, which is partially cross-cutting, since assignment to the energy manager treatment is conditional on the energy audit treatment. From the sample of 435 plants, 219 were assigned to the energy audit treatment, stratified by their electricity contract demand[11]. Only plants that completed this treatment and expressed interest in implementation were eligible for the energy manager treatment. This left an eligible group of 164 firms, of

[10] This rate appears lower per man-day than energy audits because (a) energy audits involve additional off-site analysis work (b) energy audits require the use of measurement instruments (c) the scheduling for energy managers is more flexible and hence the opportunity cost of time lower.

[11] Industrial plants declare their estimated load in advance, to help the utility forecast demand, and contract demand is the load they have signed up for with the electric utility

which 83 were assigned to the energy manager treatment.

(d) Data

Data on sample plants comes from five sources, a brief baseline survey, the energy audit reports in the treatment group, an extensive endline survey on both economic and technical outcomes, utility data on electricity consumption and expenditure, and finally a post-endline price survey of output prices using anonymous wholesale traders.

The first data source is a baseline survey covering plant characteristics such as employment and capital as well as aggregate energy use and expenditure. This survey was conducted by energy consultants and research staff together, prior to treatment assignment and coincident with the offer to enter the study sample and possibly receive free energy consulting. At the time of the baseline, plant owners or managers, who signed and stamped the survey form to register their interest in energy consultancy services.

Second, energy audit reports provide additional data on current energy consumption and projected energy savings for treatment plants. In an energy audit, consultants break down aggregate energy use into different systems in the plant and record the efficiency of these systems. For example, consultants will measure the rate of fuel input and sources of heat loss for a boiler to calculate its thermal efficiency. Energy audit reports then project, based on such calculations, what amount of energy and money would be saved were the plant to modify its operating practices or invest in new equipment. The main data of interest in energy audit reports are the recommendations, which note the system to be upgraded, its current energy consumption, the investment required for any upgrade, the projected savings and the payback period (= annual savings / investment), or number of years until the investment is projected to recoup its costs. Some recommended measures are operating or maintenance tips that carry no direct capital cost.

The third data source is a detailed endline survey covering both economic and technical aspects of each plant. The economic part of the survey comprised office interviews led by research staff with the plant owner or manager that recorded employment, materials, energy and other inputs, collected fuel and electricity bills and asked about the use of energy consultancy and recent investments around the plant in upgrades or maintenance of equipment.

The technical portion of the survey was designed to measure the efficiency of sample plants directly with physical measurements of all the main energy-using equipment found in these textile and chemical plants. This part of the survey was conducted by two energy consulting groups that employed mechanical and chemical engineers experienced in working on energy conservation with small- and medium-scale industrial plants. Thermal systems measured include the boiler, steam distribution system and process equipment, such as jet-dyeing machines or chemical reaction vessels, that are the end-users of the steam generated. Electrical systems include the plant-wide electricity distribution system as well as individual motors, air compressors and pumps that draw most of the plants' load. The survey included protocols for how to select the equipment to be measured out of the range of equipment in the plant, on the basis of fixed system characteristics or a random number table when many systems of the same type existed. Critically, this protocol did not reference the recommendations of plant energy audits in the treatment group, ensuring that the equipment selected for measurement would be comparable across the treatment and control groups of plants.

The fourth data source and primary outcome variable is electricity consumption records from the electric utilities that service sample plants. Plants were asked in the survey to give written consent for their utility to share data on their electricity consumption and expenditures. Though electricity represents only part of plant energy consumption, this administrative data provides the most accurate record of energy use available, since it is independently metered, reported on a monthly basis and available for all plants.

The fifth and final data source is a price survey conducted by wholesale textile traders to solicit output prices using an audit methodology designed to simulate a real order. Several wholesale traders were hired and given scripts for contacting select sample plants, from the textile sector only, with plant-appropriate orders for textiles in a certain quantity and with a set width, weight, color, pattern, etc. The traders recorded the price offered, after some haggling, then declined to place an order.[12]

(e) Experimental Integrity

Table 2 compares treatment and control plants using baseline survey data. Column (1) gives mean values, standard deviations, and sample sizes for each variable for treatment plants,

[12] In order to insulate traders from damage to their reputations, the plants were divided so that not too many orders would be canceled in succession and so traders would be reaching out to plants in a different area than they typically worked.

and column (2) the same statistics for the control. Column (3) reports differences estimated as the coefficient on energy audit treatment assignment in a regression of the baseline value of each variable on treatment assignment and strata fixed effects. While sample plants are mostly classed as small- and medium-enterprises, they are quite large operations. The average sample plant has 83 employees, sales of about USD 1.8 million and half a million dollars in capital.[13] Treatment and control firms are statistically balanced on these measures. Sample plants spend USD 84,000 on electricity and USD 112,000 on fuel in a year, or about 11% of sales. The audit treatment was stratified on these energy bills, so that treatment and control plants are tightly balanced on these variables at baseline. Plants use a variety of fuel sources, including lignite (low-grade brown coal, 30% of the sample) to coal (21%), diesel oil (13%) and natural gas (51%). The one significant difference noted between the treatment and control groups is that treatment plants are significantly less likely (8 percentage points on a base of 55% in the control, p-value < 0.10) to use natural gas. Fuel usage is not mutually exclusive, as plants may switch from one fuel to another.

There was some attrition in the endline survey but this is balanced across the treatment and control groups. The endline survey was conducted at least one year after the energy audit, and sometimes two years or more after the baseline survey, in order to allow time for plant to invest. Appendix Table A3 shows that 334 plants, 77% of the sample, completed the survey. A further 10% of plants had closed and 12% refused the survey, typically because the data collected was relatively more invasive than the data collected at baseline. Appendix Table A4 shows that the rate of survey completion does not significantly differ with energy audit treatment assignment. Thus the experimental sample was balanced at the time of the baseline and sample attrition does not appear to disturb this balance thereafter.

4 Results

This section presents the main results on the projected returns to energy-saving investment, actual investment and physical efficiency, plant energy consumption and finally other input and output decisions. As described in Section 2, we expect that the effects of efficiency on energy consumption would be greatest in the short-run, but may attenuate over time as other

[13] The Indian government defines small- and medium-enterprises (SMEs) as having capital stock less than INR 10 million and offers various subsidies to SMEs, which together create an incentive to understate capital investment. Employment and sales are probably more reliable measures of firm size in this context.

factors of production and output can adjust. We begin by discussing the effect of experimental treatment assignment on the completion of energy audits and the hiring of energy managers.

(a) Experimental Compliance

Assignment to treatment induced large and significant differences in the likelihood that sample plants would complete an energy audit or and smaller but significant differences in the use of an energy manager. Table 1, Panel B compares the share of plants completing energy audits and receiving additional on-site technical consultancy in the control group (column 1) for each intervention to that in the treatment group (column 2). Energy audit treatment plants are 67 percentage points (standard error 3.5 pp) more likely to receive an energy audit, a difference that is highly significant. This difference is less than 100 percent since compliance in the treatment group was imperfect, with nearly 20 percent of plants not completing audits, and some control plants got audits themselves.

Energy manager treatment plants were 23 percentage points (standard error 6.4 pp) more likely to have an energy manager, or non-audit on-site technical consultancy. While this difference is highly significant, compliance with this treatment was fairly low, with only 35 percent of the 83 plants assigned an energy manager following through. In most of these cases, the reason for non-submission of the final report was lack of interest from the plant in pursuing energy audit recommendations.

The experiment therefore generated large and significant differences in the use of energy consultancy by treatment plants, though the energy manager treatment was relatively weak and may therefore produce imprecise estimates. The results below will be reported on an intent-to-treat basis to record the effect of treatment assignment with imperfect compliance. To interpret these results as treatment-on-treated effects for plants actually using consultancy, the relevant first-stage coefficients for energy audits and energy managers are 0.67 (multiply by 1.49) and 0.23 (multiply by 4.3), respectively.

(b) Audit Projections

A candidate energy-efficiency gap is typically identified by comparing high engineering projections for returns to actual take-up or energy investment decisions (Metcalf and Hassett, 1999; Fowlie, Greenstone and Wolfram, 2015); therefore it is important to first look at what returns to energy-efficiency investments are *projected* to be. I study this by studying individual energy-saving measures recommended in energy audits and comparing the projected investments and savings for these measures to the scale of sample plants.

Figure 1 shows the distribution of investment sizes for 1,959 measures recommended for 173 treatment plants that completed energy audits. The width of each bin is USD 200. The vast majority of measures recommended require small investments below USD 1,000. Comparing investments to the scale of plants (not shown in the figure), 95% (99%) of measures require investment of less than 1.7 percent (6.6 percent) of the capital stock of the plant for which the measure was suggested.

What are the returns on these modest investments projected to be? Energy audit reports recommend investments for specific systems or pieces of equipment around the plant. Table 3 shows the characteristics of investments by the system they concern; column 1 gives the share of plants with an investment in that category, column 2 gives the number of recommendations, conditional on having any, column 3 gives the mean required investment and median annual return on investment, in terms of energy bill savings, by the type of investment.[14] The measures are ordered by the share of plants for which they are recommended. The average measure from all categories costs USD 1249, the median cost is USD 361 and the median projected return on measures recommended in 104%. The most commonly recommended measures concern lighting, motors and insulation, since these systems are present in all plants and the upgrades required are small. The median returns for these measures are projected to be in the range from 94% to 175%. The highest returns are for maintenance and other investment categories, since many of these investments have minimal investments costs, though they may have associated labor costs that consultants do not account for in energy audits.

Figure 2 summarizes the projected investments and energy savings by aggregating them across the treatment plants that completed energy audits and scaling them to the size of the average plant. All measures with any positive investment are placed in decreasing order of their projected returns. The horizontal axis then shows the cumulative investment required as a fraction of the capital stock of each sample plant to undertake these measures, and the vertical axis the cumulative projected savings in energy bills, were these measures undertaken.[15]

[14]The number of measures variable can be misleading as a sign of prevalence, since a single measure may involve the replacement of a large number of similar lights or motors.

[15]That is, $TotalInvestmentScaled = TotalInvestment/(MeanPlantCapital \times NumberOfPlants)$ and $TotalSavingsScaled = TotalSavings/(MeanPlantEnergyBill \times NumberOfPlants)$.

The figure shows that returns to energy-efficiency investment are projected to be high but diminishing. An initial investment of one percent of plant capital stock is projected to save 7.6% of plant energy bills, whereas an additional investment of one percent is projected to save an additional 2.5% of energy bills. Banerjee and Duflo (2014) use variation in the eligibility for a directed lending program to find a marginal return to capital of 105% for medium-sized Indian firms. The figure shows that if sample firms used this return as a hurdle rate, investing in all measures projected to have a higher return, they would invest a little more than two percent of capital stock and save a little more than ten percent of energy bills. Note that the extent of measures recommended, which is at the choice of consultants, does not go far beyond this level. Consultants generally report that plants will not consider investments projected to have significantly lower returns.

Energy audit reports, therefore, show that returns to energy-efficiency investment are projected to be very high but diminishing, and that the size of recommended investments is generally small. The projected quantity of investment, under a reasonable benchmark for this context, may be less than 2% of plant capital stock.

(c) Investment

This section considers actual plant investment in equipment upgrades and maintenance and physical efficiency gains in response to the experimental treatments.

Table 4 regresses actual investment on treatment status for sample plants.[16] The columns of the table show different categories of investment, in upgrades (changing a piece of equipment) or maintenance (maintaining or improving an already-existing piece of equipment). The rows show coefficients on the assignment to the two experimental treatments from the specification:

$$Investment = \beta_0 + \beta_1 EnergyAuditTreatment_i + \beta_2 EnergyManagerTreatment_i + \epsilon_i,$$

where the treatment variable represents assignment to treatment and the energy manager treatment coefficient is included in even-numbered columns.

[16] Investment in energy-using equipment maintenance and upgrades is asked of plant managers or staff with reference to specific pieces of equipment for a period of January, 2011 through the time of the endline survey, a mean time of 2.25 years.

The energy audit treatment assignment is estimated to increase investment by a modest and statistically insignificant USD 292 (column 1), which rises to USD 368 (standard error 256, p-value = 0.15 against the null of zero change in investment) when the energy manager treatment is included in the specification (column 2). This change in investment is not significant at conventional levels, though the estimated change is fairly precise and it is 48% of the base of USD 762 investment in the control. From the audit projections, we can consider an alternate hypothesis H_1 of investment equal to USD $6,981 =$ USD $10,419 \times 0.67$, which is the investment that would be undertaken with a 105% hurdle rate applied to projected returns, accounting for incomplete compliance in completion of energy audits. The table shows that this investment rule is easily rejected in the data (p-value < 0.01).

Why is the point estimate for investment negative in the energy manager treatment group? Unfortunately, due to low compliance in this group, estimates of the effect of the energy manager treatment are imprecise, so that this coefficient is not close to statistical significance and it is difficult to test alternative hypotheses regarding the effect of skill on energy consumption. Nonetheless, columns 4 and 5 attempt to provide evidence by separating investment into equipment upgrades and maintenance. The point estimate for the effect of implementation assignment on equipment upgrades is negative USD 395 (column 4) and the point estimate on equipment maintenace is positive USD 205 (column 6). Neither estimate is significantly different than zero, though in the model of column 6, testing that the coefficient on maintenance is equal to negative USD 395, the coefficient on upgrades in the model of column 4, yields a p-value of 0.12. This pattern is consistent with skilled labor, in these plants, being substitutable for energy-saving capital; plants can get by longer doing just maintenance and not upgrades if they have an engineer available.

(d) Plant efficiency

This section test for effects of the experimental treatments on plant efficiency. In the notation of the production model, the quantity of interest is $\partial f / \partial K_E$, the change in A_E with respect to additional investment.

The data collected in the endline survey provide remarkably detailed measures of plant efficiency, which are not typically available outside of a few sectors with high energy consumption and homogenous output. The primary difficulty in analysis is to construct an aggregate plant-level measure of efficiency to benchmark many possible small changes within the plant and avoid specification search. I construct an index by taking physical characteristics of each system (size, whether insulated, external temperature, etc.), identfying those that directly measure or are related to energy efficiency, standardizing these measures by substracting their mean and dividing by their standard deviation, and taking the system-level average. This creates an equally-weighted, equipment-level efficiency z-score for physical efficiency. Along with efficiency, I report results on the hours of use for each piece of equipment per day, which surveyors ask of plant staff and managers with respect to each piece of equipment.

The energy audit treatment has modest but detectable effects on plant efficiency. Table 5 relates the coefficients from regressions of efficiency indices on treatment assignment. Either the index itself (column 1) or the hours of equipment use (column 2) is the dependent variable, and the different panels record specifications for all equipment (Panel A) or for important systems in the plant, such as the boiler (Panel B), separately. Overall there is no significant increase in the efficiency index, with a coefficient of 0.042 standard deviations (standard error 0.059 standard deviations). The coefficient on the efficiency index is greater than positive 0.1 standard deviations for three of the four systems, however, with the boiler showing a positive and statistically significant 0.287 standard deviation (standard error 0.166 standard deviations) increase in the efficiency index. Notably, the hours of use for all equipment is also higher for treatment plants, by a positive and statistically significant 0.93 hours (standard error 0.45 hours) per day, on a base of 17 hours per day in the control. This roughly 5 percent increase in capacity utilization is seen across all four systems, but is greatest for the boiler system, which also had the greatest treatment effect for efficiency.

Because the efficiency index is in standardized units, it is not possible to formulate a precise elasticity of use with respect to efficiency. The evidence here does support that plants increase capacity utilization in response to efficiency.

(e) Energy consumption

Energy consumption is an important outcome for policy since the main goal of subsidized information or the promotion of energy-efficient capital is to reduce energy consumption and the externalities it causes. Though the experiment had small effects on investment and physical efficiency, it is important to consider energy consumption as an independent outcome, since the relation of energy use to investment and efficiency depends on plant responses to efficiency changes.

Electricity demand is the primary outcome for energy consumption, as explained in Section 3. Because electricity billing data is available monthly, and not only at the time of the endline survey, it is possible to estimate the effect of the treatment on electricity consumption by month relative to the timing of the treatment, like a difference-in-difference specification with randomized treatment assignment. I estimate the specification:

$$ElectricityDemand_{sitm} = \alpha_m + \alpha_s + \beta_t \times EnergyAuditTreatment_{si} + \varepsilon_{sit}.$$

The dependent variable is electricity demand (or electricity bills) for a plant i belonging to baseline electricity demand strata s in calendar year-month m (e.g., July 2013) and period t, where period is defined as $t = 0$ in the month the energy audit is conducted.[17] The strata control for pre-existing cross-sectional differences in electricity demand and the year-month dummies for seasonality and trends common to all plants. The coefficients of interest are the elements of the vector β_t, which give the treatment electricity demand relative to control in each period relative to the audit. Sample errors are clustered at the plant level to account for auto-correlation in energy consumption.

Figure 3 plots the coefficients β_t from this regression from six months prior to eighteen months after the energy audit site work. Since no period dummies are omitted, the period zero effect need not be zero by construction, but it is estimated to be very close to zero, showing the validity of the experimental design. In the year after energy audits treatment electricity consumption twice declines relative to the control group, and reaches relative lows in excess of 5,000 kWh below control consumption four and ten months after audit. However, from then onwards treatment consumption rises relative to the control and the treatment effect is estimated to be small, positive and statistically insignificant at a horizon of 18 months after energy audit.

Table 6 summarizes this time series by looking at the effect of energy audit treatment assignment on electricity demand (column 1) and bills (column 2) in the two years after audit. The first coefficient in column 1 shows that the effect of the energy audit treatment

[17] For control plants, the month relative to audit is defined according to the timing of the audit for the treatment plant in the same strata.

on monthly electricity demand is estimated to be negative 1,952 kWh (standard error 2,409 kWh), on an average monthly consumption of 56,716 kWh in the control.[18] These estimates are reasonably precise, in that one standard error represents four percent of baseline electricity consumption, whereas all measures in energy audits were projected to achieve energy savings of over 11% of energy bills if adopted. The second coefficient shows that the effect of the energy audit treatment in the second year is estimated to be of similar magnitude but opposite sign, at positive 2,249 kWh (standard error 3,871 kWh). The second column shows that these changes in electricity demand are mirrored in the electricity bill, which shows a small decline in energy charges in the first year and a small increase in the second year.

(f) Output effects

The theory predicts that, along with substituting towards greater energy consumption, plants will also reduce prices and expand output, to the extent that demand is elastic.

Since these effects are the same in the model and physical output is difficult to reliably measure, I report results on output prices. These prices were collected using wholesale textile traders as anonymous agents, in order to solicit accurate output prices for specified products from sample plants in the textile secotr. Prices in the textile sector are reported both raw, as collected, and as residuals from regressions on price on the exogenous characteristics of the mock order.

Table 7 reports results from regressions of these price measures on energy audit treatment assignment. The point estimates for both price and residual price are negative, but small and statistically insignificant. The coefficient on residual price is the larger of the two, at INR -0.9 per meter (standard error INR 1.25 per meter), a standardized effect of 0.16. Therefore I cannot reject that modest observed increases in energy efficiency did not pass through to lower output prices.

[18] As a frame of reference, an average US residential customer consumes 10,908 kWh per year, so the mean plant consumption in sample is equivalent to 62 US households.

5 Interpretation of Results: Why Does Adoption Not Respond to Information?

Despite high projected returns, sample plants invested little and saved little energy. Why was adoption of these energy-saving investments so low? I consider three different hypotheses for the difference between projections and actual take-up: that informational market failures persist, that consultants exagerrate energy savings, and that there may be ancillary fixed costs of adoption that projections do not include.

First, perhaps the treatment does not address informational market failures, for example because the energy consultants are incompetent. This interpretation of the experiment is not convincing. The consultants hired are arguably the best working in the state, and are going downmarket to work in this sample of plants: the market rate for the audit services of these consultants is about twice what control plants report paying when they get energy audits on their own. Moreover, within the sample of treatment consultants, which were randomly assigned to plants conditional on certain plant characteristics, the degree of plant investment is not significantly different across consultants. From a policy view, if the question is are their informational market failures that can feasibly be corrected, the services offered seem a strong test.

Second, consultants may have exagerrated energy savings, so that plants do not adopt. Despite considerable data on electricity consumption, it is not possible to reject the hypotheses either that energy audits achieve their projected energy savings in the first year, given the small amount of investment observed in the treatment, or that they achieve zero savings. If we assume that energy audits work *only* through investment, that plants invest in decreasing order of projected returns and that there are no changes in other inputs or capacity utilization, then the estimated USD 368 of investment is projected to yield energy savings of 1.1%. Confidence intervals for effects of the energy audit treatment on energy consumption include this projection and zero; the data is therefore consistent with this hypothesis while not positively supportive. The exagerration hypothesis alone, divorced from the adoption decision, could be tested with a (costly) technological experiment that subsidized the capital cost of adoption to measure energy savings.[19]

[19] This kind of subsidy is of policy interest. Small Industries Development Bank of India ran a USD 330m subsidized lending program for energy-efficient capital supported by the Japanese International Cooperation

Third, projections may omit ancillary fixed costs of adoption, such as plant interruption, retraining of staff or other changes to the production process. If this were the case, then the size of investments should matter as much as or more than their returns. Figure 4 shows quantile treatment effects for the energy audit treatment with investment as the dependent variable. The specification is

$$\mathcal{Q}_{Investment_i|X_i}(\tau) = \beta_1 EnergyAuditTreatment_i + \beta_2 EnergyManagerTreatment_i + \epsilon_i,$$

where $\mathcal{Q}(\tau)$ is the τ-quantile of the investment distribution conditional on treatment status. The coefficients $\beta_1(\tau)$ are reported in the figure. Most plants in both the treatment and control groups invest nothing in energy-efficiency, so quantile treatment effects up to $\tau = 0.35$ are exactly zero. Above this point in the distribution, quantile treatment effects fluctuate at a small and generally positive level until $\tau = 0.8$, at which point they turn upwards sharply. The estimated $\beta_1(\tau = 0.95) = \text{USD}72,500$ is far larger than the average treatment effects, though still narrowly statistically insignificant (p-value 0.12 against the hypothesis of no change in investment). Recall that the bulk of high-return measures recommended in energy audits were very small, so that any fixed cost could dramatically reduce their projected returns. This evidence together suggests that ancillary costs of adoption are an important possible explanation for low adoption of energy-saving measures.

6 Conclusion

This paper uses a large-scale randomized-controlled trial to study the adoption of energy-efficiency investments in a sample of industrial plants with high energy demand. Working under an existing state government subsidy program in Gujarat, India, expert energy consultants recommend adoption of high-return measures totaling two percent of capital stock to achieve savings of ten percent of energy bills. However, treatment plants receiving this advice invest only marginally more than control plants that did not.

The responses to this modest adoption appear entirely consistent with simple production theory on the response to energy-efficiency investment. Due to low adoption, the experiment is not able to sharply measure returns to energy-efficiency investment. Yet the pattern of

Agency in 2010-11.

experimental findings accords with theory: modest increases in plant efficiency and significant responses in capacity utilization of newly efficient systems, generating a temporary decline in energy consumption. Effects on output prices are of the anticipated negative sign but not statistically significant, as may be expected given modest changes in plant efficiency.

Having recorded these responses, is there an energy-efficiency gap in plants' initially low adoption? I discuss alternative hypotheses and conclude that the most powerful is that there are ancillary costs to adoption of energy-saving measures that make small but high-return investments unprofitable. Consulants account for only capital costs of investment, but plant interruptions, retraining or other may be partly fixed with respect to investment scale. These costs would lower the actual returns to investment, even if energy-saving investments indeed lower energy bills to the extent claimed.

This tepid adoption response to information on energy-saving joins a set of recent and ongoing research that casts doubt on the extent to which policy may wring savings from any energy-efficiency gap (Fowlie, Greenstone and Wolfram, 2015; Davis, Fuchs and Gertler, 2014). These findings echo, in different settings and with cleaner research designs, the findings of earlier work questioning the returns to energy-efficiency investment in the field (Joskow and Marron, 1992; Metcalf and Hassett, 1999). This research collectively is of fresh policy importance, since the U.S. EPA's Clean Power Plan, California's AB 32 and other climate policies often *assume* that energy-efficiency programs will achieve high and additional returns (Fowlie et al., 2014). If these savings do not materialize, then allowing energy-efficiency this privileged position will raise costs and undermine the efficacy of regulation.

One of the novel aspects of this paper is its focus on the efficiency of firms. The results contrast with those of Bloom et al. (2013), who find that management consulting greatly increases productivity, and profits, for Indian textile mills. Aside from differences in the domain of the consulting and measurement of outcomes, one explanation of the starkly different results may be the shape of the experimental design—intense, custom consulting for a few firms, provided by a single consultant with international reach, versus a less intensive, commercial intervention for a larger sample of plants. In the presence of fixed costs to changing plant practices and capital, the wider, shallower design of the present study would be expected to have weaker results on adoption. Indeed, the firms in Bloom et al. (2013) adopt changes rather slowly under intense consulting effort. It would be informative to observe the long-run effects of that intervention, to separate a permanent change in firm technology from a temporary change in human capital.

The last and simplest lesson for development economics is that fixed costs matter. It is difficult to accurately measure the costs of making any investment, especially when firms or consumers will change their behavior in response. Karlan, Knight and Udry (2012) find that consulting that induces micro-enterprises to adopt better, more formal accounting practices actually *lowers* profits, until firms go back to how they operated before the advice. When the stakes for the adoption of any discrete investment or technology are small, high projected returns may not reliably predict firm adoption, if they have omitted real but relatively intangible costs. Everyone wants to pick up a dollar bill on the sidewalk, but few will bend down for the infinite returns from a nickel.

References

Allcott, Hunt, Allan Collard-Wexler, and Stephen D. O'Connell. 2014. "How Do Electricity Shortages Affect Productivity? Evidence from India." *National Bureau of Economic Research Working Paper 19977.*

Allcott, Hunt, and Michael Greenstone. 2012. "Is There an Energy Efficiency Gap?" *Journal of Economic Perspectives*, 26(1): 3–28.

Anderson, S.T., and R.G. Newell. 2004. "Information programs for technology adoption: the case of energy-efficiency audits." *Resource and Energy Economics*, 26(1): 27–50.

Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2014. "Dynamic Inputs and Resource (Mis)Allocation." *Journal of Political Economy*, 122(5): 1013–1063.

Atkin, David, and Dave Donaldson. 2014. "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs." *Mimeo, UCLA.*

Banerjee, Abhijit V, and Esther Duflo. 2014. "Do firms want to borrow more? Testing credit constraints using a directed lending program." *The Review of Economic Studies*, 81(2): 572–607.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. 2013. "Does management matter? Evidence from India." *The Quarterly Journal of Economics*, 128(1): 1–51.

Davis, Lucas, Alan Fuchs, and Paul Gertler. 2014. "Cash for Coolers: Evaluating a Large-Scale Appliance Replacement Program in Mexico." *American Economic Journal: Economic Policy*, 6(4): 207–238.

Drexler, Alejandro, Greg Fischer, and Antoinette Schoar. 2014. "Keeping it Simple: Financial Literacy and Rules of Thumb." *American Economic Journal: Applied Economics*, 6(2): 1–31.

Energy Information Administration. 2013. "International Energy Outlook 2013." *Energy Information Administration, U.S. Department of Energy.*

Fischer, Greg, and Dean Karlan. 2015. "The Catch-22 of External Validity in the Context of Constraints to Firm Growth." *Forthcoming, American Economic Review, Papers & Proceedings.*

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98(1): 394–425.

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. 2015. "Are the Non-Monetary Costs of Energy Efficiency Investments Large? Understanding Low Take-up of a Free Energy Efficiency Program." *Energy Institute at Haas Working Paper*, 256.

Fowlie, M, L Goulder, M Kotchen, S Borenstein, J Bushnell, L Davis, M Greenstone, C Kolstad, C Knittel, R Stavins, et al. 2014. "Environmental economics. An economic perspective on the EPA's Clean Power Plan." *Science (New York, NY)*, 346(6211): 815–816.

Hausman, J.A. 1979. "Individual discount rates and the purchase and utilization of energy-using durables." *The Bell Journal of Economics*, 10(1): 33–54.

Hsieh, Chang-Tai, and Peter J Klenow. 2009. "Misallocation and manufacturing TFP in China and India." *The Quarterly Journal of Economics*, 124(4): 1403–1448.

Hsieh, Chang-Tai, and Peter J Klenow. 2014. "The Life Cycle of Plants in India and Mexico." *The Quarterly Journal of Economics*, 129(3): 1035–1084.

Jacobsen, Grant D., and Matthew J. Kotchen. 2010. "Are Building Codes Effective at Saving Energy? Evidence from Residential Billing Data in Florida." , (16194).

Joskow, P.L., and D.B. Marron. 1992. "What Does a Negawatt Really Cost? Evidence from Utility Conservation Programs." *The Energy Journal*, 13(4): 41–74.

Karlan, Dean, Ryan Knight, and Christopher Udry. 2012. "Hoping to Win, Expected to Lose: Theory and Lessons on Micro Enterprise Development." *NBER Working Paper 18325.*

Metcalf, G.E., and K.A. Hassett. 1999. "Measuring the energy savings from home improvement investments: evidence from monthly billing data." *Review of economics and statistics*, 81(3): 516–528.

Rohdin, P., P. Thollander, and P. Solding. 2007. "Barriers to and drivers for energy efficiency in the Swedish foundry industry." *Energy Policy*, 35(1): 672–677.

Syverson, Chad. 2004. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy*, 112(6).

Train, Kenneth. 1985. "Discount rates in consumers' energy-related decisions: a review of the literature." *Energy*, 10(12): 1243–1253.

7 Figures

Figure 1: Capital Needed for Energy-Efficiency Investments

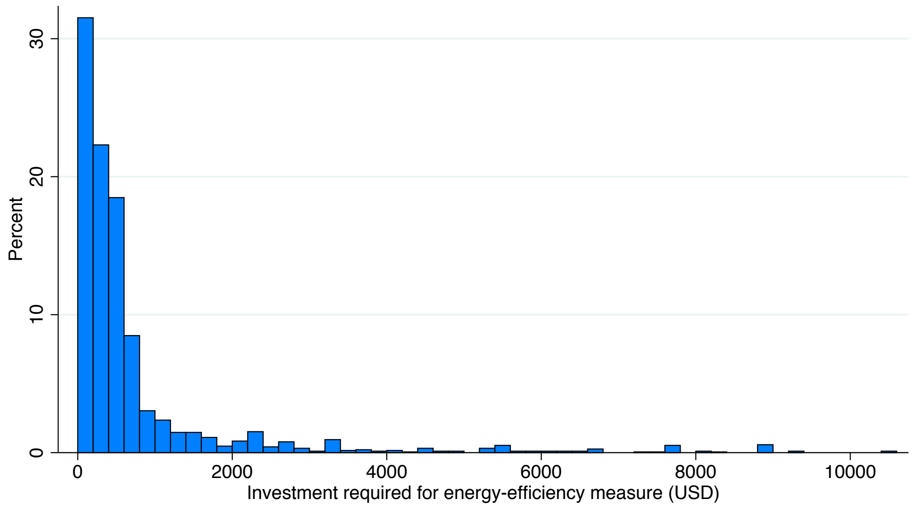

The figure shows the distribution of investment costs for measures recommended in energy audits in the energy audit treatment group of plants. Each bin is USD 200 wide and the distribution is truncated at the 97.5th percentile (USD 11,111) for clarity. The investments are for 1,959 different measures recommended to 173 treatment plants.

Figure 2: Projected Returns on Investment in Energy-EFficiency

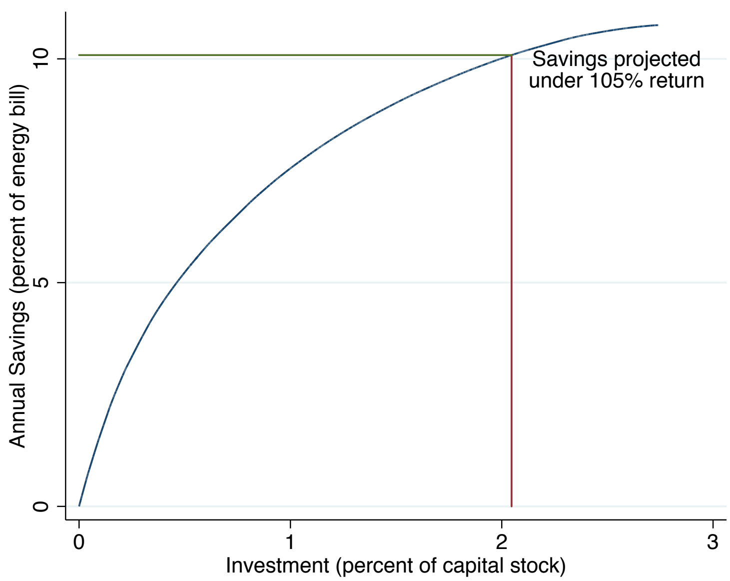

The figure summarizes the projected investments and energy savings in energy audits of treatment plants by putting measures in decreasing order of projected returns and then scaling cumulative investment and savings as a fraction of the size of the average plant. That is, $TotalInvestmentScaled = TotalInvestment/TotalPlantCapital$ and $TotalSavingsScaled = TotalSavings/TotalPlantEnergyBill$. The vertical and horizontal lines indicate the investments and savings projected if firms applied a hurdle rate of a 105% annual return to their investment decisions, where this rate of return is drawn from the estimate of marginal returns to capital for Indian firms in Banerjee and Duflo (2014).

Figure 3: Electricity Use Over Time: All Plants

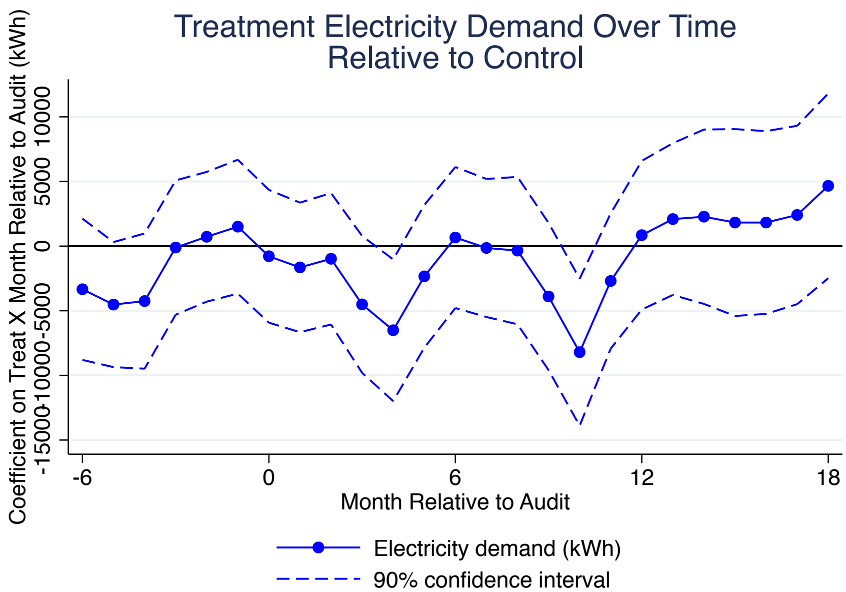

The figure shows regression coeficients estimating the effect of the energy audit treatment on electricity consumption by month relative to the timing of the energy audits for treatment plants. The specification includes strata fixed effects for baseline levels of electricity consumption and year-month fixed effects (e.g., July 2013) to control for seasonality and trends common to all plants. The month relative to audit is defined as $t = 0$ in the month the energy audit is conducted.

Figure 4: Energy Audit Assignment Quantile Treatment Effects

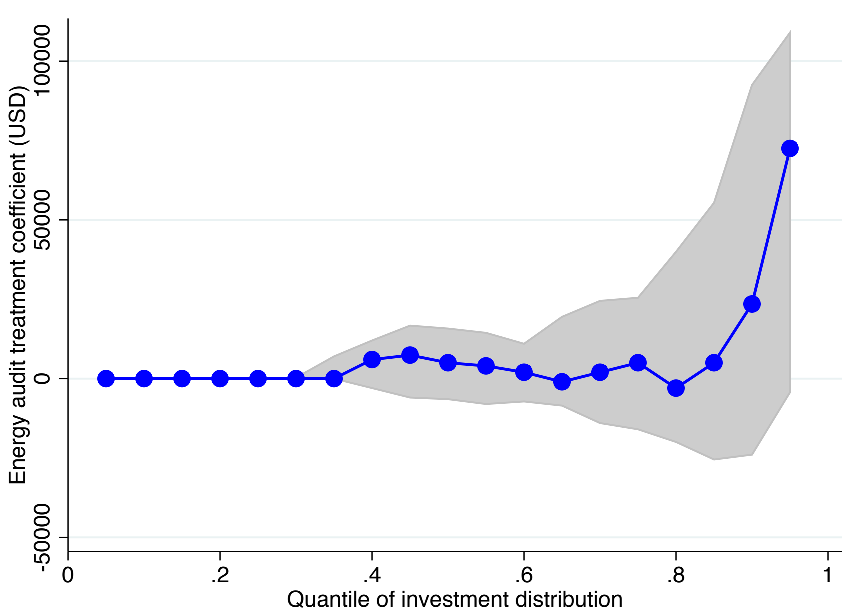

The figure shows quantile treatment effects for the energy audit treatment with investment as the dependent variable. Coefficients $\beta_1(\tau)$ are reported for quantiles from $\tau = 0.05$ to $\tau = 0.95$ in increments of 0.05. All specifications control for energy manager treatment assignment.

8 Tables

Table 1: Experimental Design and Compliance

	Control (1)	Treatment (2)	Total (3)
A. Treatment Assignments			
	Energy Audit Treatment		
Energy Manager Treatment			
Total	216	219	435
Control	0	81	81
Treatment	0	83	83
Not assigned	216	55	271
B. Treatment Completion			
Energy Audit Treatment			
Share of plants doing energy audit	0.12	0.79	0.67***
			(0.035)
Energy Manager Treatment			
Share of plants using energy manager	0.12	0.35	0.23***
			(0.064)

The table shows the experimental design and treatment assignments, in Panel A, and the actual completion of energy audits or use of energy managers, in Panel B. In Panel A the columns indicate energy audit treatment assignment and the rows energy manager treatment assignment and each cell reports the number of plants assigned to that combination of treatments. Plants are assigned to the energy manager treatment conditional on completing the energy audit treatment and expressing interest in implementation. In Panel B the columns indicate treatment assignment status for the energy audit treatment, in the first row, and the energy manager treatment, in the second row. Each table entry gives the share of plants completing the treatment, or similar consultancy outside the experiment, conditional on the treatment assignment shown by the column.

Table 2: Balance of Baseline Covariates by Energy Audit Treatment

	Treatment (1)	Control (2)	Difference (3)
Contract demand (kVA)	200.9	191.9	8.98
	[172.1]	[171.7]	(16.5)
Electricity bill (Annual USD 000s)	85.7	82.6	3.08
	[109.9]	[106.3]	(10.4)
Fuel bill (Annual USD 000s)	110.2	114.9	-4.63
	[428.5]	[275.0]	(34.6)
Employees	83.6	82.7	0.97
	[112.7]	[117.5]	(11.2)
Capital (USD 000s)	529.0	581.6	-52.6
	[750.3]	[813.1]	(81.4)
Sales (USD 000s)	1677.2	1809.9	-132.7
	[2427.7]	[3725.4]	(320.2)
Uses lignite (=1)	0.29	0.32	-0.029
	[0.46]	[0.47]	(0.045)
Uses coal (=1)	0.23	0.19	0.036
	[0.42]	[0.39]	(0.039)
Uses diesel oil (=1)	0.11	0.16	-0.051
	[0.31]	[0.37]	(0.032)
Uses gas (=1)	0.47	0.55	-0.081*
	[0.50]	[0.50]	(0.048)
Observations	217	216	433

The table shows means of baseline characteristics for treatment and control plants and the difference between these groups, estimated as the coefficient on energy audit treatment assignment in a regression of each outcome on treatment and a set of baseline electricity use strata dummies. Standard deviations of each variable are in brackets and standard errors of each estimated difference in parentheses. Employment, capital and sales are reported at baseline for 422, 369 and 383 total plants, respectively. Statistical significance of differences is marked by * $p < 0.10$, ** $p < 0.05$ and *** $p < 0.01$.

Table 3: Frequency of Energy-Saving Measures

	Measure Prevalence		Investment Size (USD)		Return (%)
	Plant Share (1)	Number if Any (2)	Mean (3)	Median (4)	Median (5)
Total	1.00	12.04	1249.06	361.11	104
Lighting	0.82	1.57	1304.40	305.56	94
Motor sizing / efficiency	0.78	10.13	1029.29	361.11	101
Insulation	0.46	2.17	476.50	260.00	175
Electricity Tariff	0.42	1.40	579.40	253.33	154
Heat Recovery	0.39	1.27	5956.81	5581.51	220
Maintenance / Other	0.23	1.38	1083.40	536.11	1601
Automation	0.10	1.06	2017.35	666.67	155
Compressors	0.09	1.20	3495.29	1111.11	80
Drives / belts / pulleys	0.07	1.91	2666.51	2000.00	71

The table shows characteristics of investments recommended in energy audits of treatment plants. The rows give the type of equipment or system that the measure involves, and the columns give statistics on the prevalence (columns 1 and 2), investment cost (columns 3 and 4) and returns (column 5) on these measures. A total of 1,959 measures have non-zero investment and so finite returns.

Table 4: Equipment-Level Investment on Treatment

	Total Investment (USD)		Upgrades (USD)		Maintenance (USD)	
	(1)	(2)	(3)	(4)	(5)	(6)
Energy audit treatment (=1)	291.8	367.7	106.9	264.9	184.9	102.8
	(282.1)	(256.3)	(193.9)	(220.4)	(183.9)	(129.8)
Energy manager treatment (=1)		-189.6		-394.5		204.9
		(560.6)		(358.2)		(384.5)
Control mean	762.2	762.2	323.8	323.8	438.4	438.4
p-value for H_0 : Energy audit treatment = 0	0.30	0.15	0.58	0.23	0.32	0.43
p-value for H_1 : Energy audit treatment = 6981	0.00	0.00				
Observations	329	329	329	329	329	329

The table shows coefficients from regressions of investment in energy-efficient equipment at the plant level on treatment status for sample plants. Investment in energy-using equipment maintenance and upgrades is asked of plant managers or staff in the Endline Survey with reference to specific pieces of equipment for a period of January, 2011 through the time of the endline survey, a mean time of 2.25 years. The columns of the table show different categories of investment, in upgrades (changing a piece of equipment) or maintenance (maintaining or improving an already-existing piece of equipment). The rows show coefficients on the assignment to the two experimental treatments from the specification, with standard errors in parentheses beneath and statistical significance indicated by * $p < 0.10$, ** $p < 0.05$ and *** $p < 0.01$. Regressions include fixed effects for baseline electricity consumption strata.

Table 5: Physical Efficiency of Plant Equipment on Treatment Status

	Efficiency (1)	Hours used per day (2)
Panel A. All Systems		
Energy audit treatment (=1)	0.0421	0.931**
	(0.0593)	(0.451)
Control mean	0.0	16.7
Observations	2175	2134
Panel B. Boiler		
Energy audit treatment (=1)	0.287*	1.922*
	(0.166)	(1.142)
Control mean	-0.1	16.3
Observations	292	283
Panel C. Motors		
Energy audit treatment (=1)	-0.0129	0.746
	(0.0786)	(0.497)
Control mean	0.0	17.0
Observations	1570	1544
Panel D. Jet-Dyeing Machines (Textile Plants Only)		
Energy audit treatment (=1)	0.142	1.100
	(0.250)	(1.157)
Control mean	-0.0	19.4
Observations	128	128
Panel E. Reaction Vessels (Chemical Plants Only)		
Energy audit treatment (=1)	0.152	0.965
	(0.360)	(2.256)
Control mean	-0.0	13.5
Observations	185	179

The table shows regressions of the physical efficiency and hours of use of pieces of equipment within plants on energy audit treatement status. Physical efficiency is measured with a standardized index of efficiency composed of the average of standardized physical efficiency measures, such as the presence of insulation or the external temperature of an insulated vessel, where each standardized measure is signed so that efficiency is monotonically increasing in the index. These measures are recorded at the level of the piece of equipment and so multiple measures may be available per plant. Hours of operation are obtained by asking plant managers or staff how often different equipment is run. Regressions include baseline electricity demand strata and controls for energy manager treatment assignment. Standard errors, clustered at the plant level, are in parentheses, with statistical significance marked by * $p < 0.10$, ** $p < 0.05$ and *** $p < 0.01$.

Table 6: Electricity Demand on Treatment

	(1) Demand (kWh)	(2) Bill (USD)
Energy audit treatment (=1) × First year post audit	-1951.5	-67.70
	(2409.2)	(253.2)
Energy audit treatment (=1) × Second year post audit	2348.9	98.96
	(3870.6)	(418.0)
Control mean	56715.8	7164.8
Plants	322	322
Observations Per Plant	24.3	24.3
Observations	7811	7809

The table shows coefficients from regressions of monthly electricity demand and energy charges on interactions between energy audit treatment assignment and the time since energy audit in months. The regressions include controls for baseline electricity demand strata and year-month fixed effects (e.g., July 2013). Standard errors clustered at the plant level in parentheses with statistical significance indicated by * $p < 0.10$, ** $p < 0.05$ and *** $p < 0.01$.

Table 7: Output Prices on Treatment

	(1) Price (INR/meter)	(2) Price residual (INR/meter)
Audit treatment assignment (=1)	-0.0685	-0.901
	(1.111)	(1.251)
Control mean	12.6	0.8
Control sd	7.0	5.7
Observations	276	255
Plants	67	62
Observations Per Plant	4.1	4.1

The table shows coefficients from regressions of output prices on energy audit treatment assignment, where prices are solicited by working textile traders from select sample plants in the textile sector using an audit methodology. The dependent variable in the first column is raw price data and in the second column the residual from a regression of price on the characteristics of the order placed by the trader, which are exogenously assigned. The regressions include controls for baseline electricity demand strata. Prices are recorded for multiple products for each plant. Standard errors clustered at the plant level in parentheses with statistical significance indicated by * $p < 0.10$, ** $p < 0.05$ and *** $p < 0.01$.

A Appendix

Table A1: Plant Interest in Energy Audit

	Plant interest	
	Number	Percent
	(1)	(2)
Interested	490	53.0
Already have consultant	38	4.11
Energy not a large cost	40	4.32
Scope of savings not large	50	5.41
Other	307	33.2
Total	925	100

The table shows interest in joining the experimental sample, as solicited by energy consultants, amongst plants in the larger population of membership rolls in local industrial associations.

Table A2: Selection into the Sample

	Sample mean [sd]		
	Interested	Not	Difference
Electricity load in kW	65.9	60.4	5.53
	[67.4]	[78.0]	(11.3)
Capital (USD '000s)	231.7	130.6	101.0
	[445.0]	[165.5]	(63.2)
Capital, plant (USD '000s)	96.5	77.0	19.5
	[184.8]	[111.1]	(27.2)
Capital, equipment (USD '000s)	7.72	7.39	0.33
	[41.0]	[36.5]	(6.41)
Capital, building (USD '000s)	76.8	29.1	47.6
	[219.3]	[40.6]	(30.6)
Capital, land (USD '000s)	50.6	17.2	33.5
	[246.4]	[14.8]	(34.3)
Employment	15.2	13.2	1.91
	[13.9]	[11.4]	(2.14)
Observations	154	52	

The table compares observable characteristics of plants interested or not interested in participating the experiment in data from industrial registrations with the state of Gujarat. Industrial registration data is available for a subset of 206 of the 925 plant population fo the experiment due to partial registration and limited matching. Registration data gives characteristics as reported to the government on establishment of the plant.

Table A3: Attrition in the Endline Survey

	N	%
Surveyed	334	76.8
Not surveyed	101	23.2
Shut down/Sold off	42	9.7
Refused	52	12.0
Other	7	1.6
Total	435	100.0

Shut down/sold off includes plants that were permanently closed and plants that were temporarily closed during repeated survey visits. *Refused* includes plants that were operating at the time of the visit, but that refused to respond. *Other* includes plants that moved or that could not be contacted.

Table A4: Endline Attrition by Treatment Status

	Treatment	Control	Difference
Surveyed	0.744	0.792	-0.048
	[0.437]	[0.407]	(0.040)
Shut down/Sold off	0.096	0.097	-0.001
	[0.295]	[0.297]	(0.028)
Refused	0.142	0.097	0.045
	[0.349]	[0.297]	(0.031)
Other	0.018	0.014	0.004
	[0.134]	[0.117]	(0.012)
Observations	219	216	

Shut down/sold off includes plants that were permanently closed and plants that were temporarily closed during repeated survey visits. *Refused* includes plants that were operating at the time of the visit, but that refused to respond. *Other* includes plants that moved or that could not be contacted.