NBER WORKING PAPER SERIES

MEASURING THE ENERGY SAVINGS FROM HOME IMPROVEMENT INVESTMENTS: EVIDENCE FROM MONTHLY BILLING DATA

Gilbert E. Metcalf
Kevin A. Hassett

Working Paper 6074

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 1997

This project was supported in part by funding from the National Science Foundation (SBR-9514989) to the NBER and by the Department of Energy. We are grateful for comments from Robert Pindyck and for expert data assistance from Chris Myers and Lara Shore-Sheppard. This paper is part of NBER's research program in Public Economics. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Board of Governors of the Federal Reserve System or the National Bureau of Economic Research.

Measuring the Energy Savings from
Home Improvement Investments:
Evidence from Monthly Billing Data
Gilbert E. Metcalf and Kevin A. Hassett
NBER Working Paper No. 6074
June 1997
JEL Nos.E22, Q40, Q48
Public Economics

ABSTRACT

An important factor driving energy policy over the past two decades has been the "Energy Paradox," the perception that consumers apply unreasonably high hurdle rates to energy saving investments. We explore one possible explanation for this apparent puzzle: that realized returns fall short of the returns promised by engineers and product manufacturers. Using a unique data set, we find that the realized return to attic insulation is statistically significant, but the median estimate (12.3 percent) is close to a discount rate for this investment implied by a CAPM analysis. We conclude that the case for the Energy Paradox is weaker than has previously been believed.

Gilbert E. Metcalf
Department of Economics
Tufts University
Medford, MA 02155
and NBER
gmetcalf@tufts.edu

Kevin A. Hassett
Board of Governors of the
Federal Reserve System
Mailstop 80
Washington, DC 20551

I. Introduction.

Home owners have available to them a wide variety of energy saving home improvements, many of which can--according to engineering estimates--pay for themselves in a very short time period. One of the main puzzles facing policy makers over the past two decades has been the persistent observation that home owners appear to apply very high discount rates to these investment opportunities (see, e.g. Hausman (1978)). What we have elsewhere dubbed the Energy Paradox (see Hassett and Metcalf (1994)) -- the apparent use of extremely high discount rates for home improvement investments -- can perhaps be explained by drawing on insights from the literature on irreversible investment (see Dixit and Pindyck (1994) for an excellent treatment of this literature). If energy conservation investments have low salvage value and their return is risky, then optimal investment hurdle rates can exceed the underlying discount rate by a substantial margin (Metcalf and Rosenthal (1995) develop this argument in the context of energy efficient lighting and refrigerators).

In this paper, we explore an alternative hypothesis: namely, we investigate whether discount rates in past studies were overestimated because returns to energy saving investments were overestimated. When past researchers have estimated consumer discount rates, they have used the engineering rates of return as a key input. If an optimizing person who is not liquidity constrained refused to invest in a project with a 20 percent rate of return, for example, then we can infer that her hurdle rate for this type of investment is bounded below by 20%. One possible explanation for the high estimated discount rates in past studies is that consumers do not expect to receive anything like the engineering return if they were to adopt the home improvement. This might be the case if, for example, engineering estimates of potential energy savings--which are often provided by the manufacturer of the relevant product--misrepresent savings because they are based on highly controlled studies which do not directly apply to actual realized savings in a representative house.

We analyze whether realized returns are comparable to technical return estimates using the RECS (Residential Energy Consumption Survey) compiled by the U.S. Department of Energy. The RECS survey carefully records household investment in many different energy conserving devices for a cross section of households, and for this study we focus on the returns to attic insulation. We chose to do this because attic insulation is measured very well in our data, and the potential benefits to it are---according to controlled experiments, at least---large enough that one would expect any reasonable methodology to detect them. In addition, the survey collects monthly energy bill data to monitor energy consumption. We find support for the view that realized returns are smaller than has been suggested by past research, with fully 50 percent of our sample likely receiving less than a 13 percent annual return to their insulation investment.

II. Review of the Literature

While virtually every home improvement product on the market comes with an engineering based estimate of the potential energy savings from the use of that product, there have been relatively few studies documenting the actual returns received by individuals who pursue home improvement strategies. In a series of papers that are closest in spirit to our work, Hirst (1987) and Hirst and Goeltz (1984,1985) evaluate the effects on energy use of the Bonneville Power Administrations interim residential weatherization program. This program performed energy audits--from 1980-82--free to customers, and included zero interest loans for installation of measures recommended by the audits. These researchers study programs which focused on houses that used electricity for heating, and estimated the energy savings for customers that performed retrofits. According to their estimates, the energy savings of these programs were sufficient to cover the costs of the retrofits, which averaged $2650 per household in their sample. Train (1985) reports hurdle rate measures for thermal integrity based on national surveys of 26% and 32% based on attic insulation.

Hartman and Doane (1987) studied the Portland General Electric Companies audit program, and found that consumption dropped for participants in the program. Sebold and Fox (1985) studied the San Diego Gas and Electric Company's audit program, and compared the realized returns by individuals to that predicted by engineering studies. They found that, on average, total returns came in somewhat below those predicted by engineering studies. A similar finding is reported in Hirst (1986) who finds that actual savings from retrofit programs fall short of savings predicted by energy auditors by 22 to 53%. The Sebold and Fox results along with the Hirst results lend support for our hypothesis that engineering or other "professional" estimates of returns are biased upwards.

Our study differs from the past work both in its methodology and its focus. First, rather than investigating the total return to a comprehensive, utility sponsored project, we analyze the returns for individual projects that consumers undertake. We think that this focus is important for a number of reasons. First, participants in a utility program presumably receive careful coaching about energy use from the auditors. It may be that the energy savings that follow the investment could be received without making any energy investments.[1] Since our sample, which

[1] For example, participants may be more conscious of energy waste, and may be quicker to shut doors and windows during heating and cooling season.

is discussed in more detail below, does not include only improvements by people involved in audits, we will be able to more closely approximate the actual savings received by typical individuals who pursue home improvement without careful coaching. Second, by focusing on individual projects, we can construct measures of the actual returns achieved in a typical home. This is an important new step, since engineering estimates provided by manufacturers of potential returns may not accurately reflect the true returns that households could expect to receive.[2] In addition to shedding light on consumer behavior, our estimation methodology may be helpful in formulating more efficient energy conservation plans in the future.

III. Measuring the Returns to Home Improvement Investments

In this section we present the basic model we use to estimate the energy savings experienced by households in our sample. One could measure the returns to conservation investment by regressing energy consumption on attributes of the family and the house, and dummy variables indicating investments made in energy efficient capital. The coefficient on the investment dummy would provide a measure of the returns to the investment (in terms of energy reductions). The problem with the regression as described is that weather conditions vary dramatically across households and across time. Hence, if a house puts in a high efficiency fuel burner and we observe a decrease in energy consumption, we cannot identify whether the

---

[2] Even if insulation firms are not artificially inflating estimates, this might happen if the assumptions used to generate predicted returns do not accurately reflect the properties of the capital stock that is being modified. For example, an estimate of the return to weatherstripping would also depend on the thermal integrity of the window that is being treated. (If the window were open, then the return would be zero). When calculating an engineering return, some assumption about the thermal integrity of the treated window must be made. It may be that the characteristics of households that engage in home improvement investments differs significantly from those assumed in baseline engineering savings calculations.

conservation investment lowered fuel use or whether winter temperatures happened to moderate in the second year. One approach to this problem would be to include temperature data on the right hand side - both alone and interacted with key variables. However, this is a somewhat ad hoc weather adjustment. Fortunately, there is an adjustment we can make that follows from energy engineering considerations. Rather than use annual fuel consumption, we use Normalized Annual Consumption (NAC), a measure based on a normalization approach used in the Princeton Scorekeeping Method (PRISM), developed by the Center for Energy and Environmental Studies at Princeton University. We begin with a brief description of the construction of this measure.[3]

The demand for space heating energy (E_h) follows from engineering principles. The space heating energy required to maintain a desired indoor temperature of T_i is proportional to the difference between T_i and outdoor temperature (T_o):

(1) $$E_h = \lambda(T_i - T_o)$$

where λ measures the inherent "tightness" of a house. For example, a house with significant air infiltration losses would have a higher value of λ. Energy heating is supplied by fuel (ϕ) which is burned with efficiency η and by heat loss from appliances, people, and solar gain (Q):

(2) $$E_h = \eta\phi + Q$$

Both the efficiency factor η and solar gain (Q) are functions of housing characteristics. For example, a house with new double paned glass would increase Q without increasing λ. Equating (1) and (2), fuel requirements are given by:

(3) $$\phi = \lambda(T_i - T_o)/\eta - Q/\eta$$

or

[3] The PRISM methodology is described in Fels (1986).

(4.1) $$\phi = \beta(\tau - T_o)$$

(4.2) $$\beta = \lambda/\eta \text{ and}$$

(4.3) $$\tau = T_i - Q/\lambda.$$

The variable τ can be thought of as the desired temperature setting adjusted for natural heating from other sources (Q). Since fuel consumption is restricted to be non-negative, equation 4.1 can be rewritten as

(5) $$\phi = \beta \cdot HDD(\tau)$$

where $HDD(\tau) = \max(0, \tau - T_o)$ is the instantaneous heating degree day measure for the house at reference temperature τ.[4] Finally if the house uses a constant amount of fuel for other purposes (lighting, washing, etc.) at rate α, fuel use becomes

(6) $$\phi_{it} = \alpha_i + \beta_i \cdot HDD(\tau)_{it} + \epsilon_{it}$$

where i indexes houses and t indexes billing periods and an error term has been added the equation.

Normalized Annual Consumption (NAC) in any given year can then be constructed as

$$NAC = 365\hat{\alpha} + \hat{\beta} \cdot HDD_N(\tau). \qquad (7)$$

where $HDD_N(\tau)$ is the amount of heating degree days in a typical year for a house in this region and $\hat{\alpha}$ and $\hat{\beta}$ are house specific estimates from a regression of energy consumption on heating degree days from equation (6). HDD_N is computed as the average heating degree at base τ over an historical period. The use of normalized annual consumption rather than actual consumption

[4] This is an approximation as conventionally measured heating degree days do not adjust for natural heating within a home.

for the dependent variable purges the energy consumption data of short term weather fluctuations.

The NAC measure in (7) assumes that energy is only used to heat homes. In many parts of the country, a more significant use of energy is for home cooling. An analogous derivation to the one leading to equations (6) and (7) leads to our adding a measure of cooling degree days (CDD) to control for air conditioning:

(8) $$\phi_{it} = \alpha_i + \beta_{1i} \cdot HDD(\tau)_{it} + \beta_{2i} \cdot CDD(\tau)_{it} + \epsilon_{it}$$

and

$$NAC = 365\hat{\alpha} + \hat{\beta}_1 \cdot HDD_N(\tau) + \hat{\beta}_2 \cdot CDD_N(\tau') \tag{9}$$

Since housing characteristics and taste towards heating and cooling may vary widely, it is important to allow as much variation in α and ß as possible. In this study, we estimate α and ß at the individual house level, by regressing average daily consumption in month t on average heating degree days $(HDD/N)_t$ for the month. Monthly consumption data and heating degree data are available for roughly 500 houses in the RECS data set in 1984 and 1990, and for about 1000 households in 1987. For each year, we run k regressions, where k is the number of households in our study, each on roughly 12 observations,[5] and construct house specific estimates of α and ß in each year. Hence we generate estimates of NAC_{kt}, k indexing houses and t indexing years (1984, 1987, and 1990).

Once we have calculated the NAC for each household, calculating the savings from different types of home improvement investment will simply involve attempting to predict

---

[5] The unit of observation is actually billing periods rather than months. The typical billing period is a month.

variation in the NAC using--among other things--structural characteristics of the homes we study. Since we have panel data on households, and the energy billing data for individual households both before and after home improvement investments are adopted, we should have a great deal of power to identify energy savings--should they exist--associated with even fairly small scale home improvements. We thus proceed to regress consumption on characteristics of the household, the housing structure including measures of the presence of energy conservation investments:

$$\ln(NAC_{kt}) = \alpha_k + \beta X_{kt} + \epsilon_{kt} \tag{10}$$

The vector X includes family characteristics (including number of children and household income for example), and housing characteristics (e.g. house size). It also includes dummy variables for the presence of various conservation investments. The coefficient on the dummy variable for the presence of the investment gives the percentage change in energy consumption holding other characteristics of the house and family constant. In this study, we focus exclusively on attic insulation. This is a popular investment that can be done by a handy homeowner and has the potential for substantial returns. Unlike some other investments tracked in the RECS data, it is fairly easy to measure the extent of this form of investment.

No matter how complete the listing of housing and individual characteristics is, it may be that such a regression does not capture adequately all variables that affect the energy consumption of households. In this case, omitted variable bias may significantly alter our empirical results. To some extent, however, it may be reasonable to presume that the unobserved variables are constant across households, and hence, a fixed effects estimator may provide a better estimate of the marginal impact on energy consumption of a home improvement. As a final step, we explore this possibility.

To summarize, we proceed as follows:

1) We create monthly data on temperature and energy consumption for each household in each year.

2) We estimate equation (8) for each household in each year, in order to obtain individual estimates of both α and β.

3) We generate estimates of NAC for each household using household specific coefficients.

4) We regress NAC on energy capital, housing and demographic variables in order to identify the energy savings associated with specific characteristics.

IV. The Data

Our main data source is the Department of Energy's Residential Energy Conservation Survey (RECS).[6] The RECS survey is conducted in two major parts: the Household Survey and the Energy Supplier Survey. In the household survey, information concerning the structure in which the household lives and demographic information is gathered. The Energy Supplier Survey contains billing records of actual energy consumption for the surveyed households.[7]

We have available to us three waves of RECS data: 1984, 1987 and 1990. The survey is designed to collect information about households in two successive surveys. Hence, in each wave, half the households are held over from the previous survey while half are added (to be

---

[6] Results of the RECS have been reported by the Office of Energy Markets and End Use in a series of publications.

[7] An important attribute of the RECS survey is the acquisition of energy consumption data from energy suppliers directly. This eliminates a major source of noise in the data resulting from relying on householder's recollections of past energy consumption.

reinterviewed in the next survey). This rotation design allows us to construct a panel of households observed in two periods.

Table 1 summarizes the coverage of our sample. Each survey comprises between five and six thousand households. However, we lose a number of households for various reasons. First, we focus on households who list their primary heating source as Gas or Electric. We limit ourselves to these households because monthly billing data in 1987 are only available for these households. This reduces our sample by roughly a quarter. Next we throw out households with insufficient monthly billing data. For purposes of running the individual regressions in equation (8), we impose the restriction that there be at least 12 observations on billing data for the household specific regression. We also eliminate households with excessively long billing periods (defined as 70 or more days in length). This reduces measurement error resulting from averaging temperature and consumption data over long time periods.[8] These restrictions reduce our sample by 44%. Finally, we consider only those households present in both years of the sample. This gives up 2272 observations on 1136 households.

We occasionally will restrict our sample further. First, we eliminate households with very large changes (50% change in absolute value) in the measured normalized annual consumption measure from the household level regressions from the first year to the next. This restriction reduced the sample by 5%. While these outliers appear to result from poorly estimated first stage regressions (NAC regressions), it is a judgement call whether to drop them or not. We report regressions with and without these outliers; as will be seen, the results are not sensitive to dropping them. We also lose households in regressions in which we enter household specific

[8] A few households have billing periods as long as 300 or more days.

information. Some of the variables we use in the second round regressions are missing for some households. The largest loss is due to missing data on age of furnace equipment. This reduces the number of households by 300, from 1136 to 836. We also find that the type of housing unit makes a big difference in the second round regressions. The survey looks at single family homes, both detached and attached, housing units in large buildings (typically rental units in apartment houses), and mobile homes. Restricting analysis single family units reduces the sample of households to 945. Imposing all three restrictions reduces the sample from 1136 households to 765.

V. Normalized Annual Consumption

Table 2 summarizes our estimates of the coefficients in equation 8. Recall that these are house-year specific regressions of energy consumption on heating and cooling degree days. Variation in these coefficient estimates reflects variation in housing structure as well as individual behavior. This variation will be reflected in variation in our measured NAC to be constructed below. Each element of the first row for each variable in the table gives the change in million BTUs per change in hundred degree days. Standard deviations are in parentheses. For example, our mean estimate for 1984 suggests that an extra hundred heating degree days (roughly 3 degrees colder each day for the entire month) would increase energy consumption in that month by 1.617 million BTUs[9]. There is wide dispersion in the responsiveness of energy consumption to temperature swings, however. The next two numbers--rows 3 and 4 under each variable, give the 25th and 75th percentile values for the estimates.

---

[9] There are roughly 7.8 gallons of fuel oil per million BTUs. Thus, this increase in energy consumption is about 12 ½ gallons of fuel oil.

We next proceed to constructing estimates of normalized energy consumption (NAC). We use the coefficient estimates from equation 8 for each household in each year to predict consumption assuming "normal" weather conditions. Normal weather conditions are defined as average heating and cooling degree days in the sample for each household.[10] Plugging these data into equation 9 yields our measure of NAC. Table 3 provides summary information on normalized annual consumption. Normalized energy consumption ranges widely in the sample, with the 75th percentile being roughly double the 25th percentile. Wave 1 refers to households that entered the survey in 1984 (exiting in 1987) and wave 2 refers to households entering survey in 1987 (exiting in 1990). Focusing on 1987, it appears that the wave 1 households have slightly higher energy consumption (controlling for weather). We now proceed to the second stage of analysis.

VI. Estimating the Energy Savings of Investments

To estimate the effects of energy conservation investments on energy consumption we estimate a semi-log regression, where the log of each individual's NAC is regressed on conservation and control characteristics. Given this specification, the coefficient on each variable is interpretable as the percent change in energy use due to a change in that variable. We begin by reporting sample statistics on the data (Table 4). Houses in the sample tend to be old (with half of them built in the 1950s or earlier) and with old heating systems. Roughly 3/4 of the houses have attic insulation (with some adding insulation during the period of observation). We also include information on the characteristics of the house (number of windows, area heated), heating and

[10] A better measure would use average heating and cooling degree days over a longer time period. Unfortunately, location identifiers are not provided in the RECS data preventing us from constructing these measures.

cooling practices (thermostat setting and use of air conditioners) as well as price information for gas and electricity. Price1 is the price of electricity for households using electricity as their main heating source, price2 is the price of gas for households using gas as their main heating source, and price 3 is the price of electricity for households using gas as their main heating source. Each of these prices are in cents per 1000 BTUs.

Our basic regression is of the form

(10') $$\ln(NAC)_{it} = \alpha_i + \beta_1 K_{it} + \beta_2 X_{it} + \epsilon_{it}$$

where K is a dummy variable indicating the presence of attic insulation, and X is a vector of variables measuring characteristics of the house and the family. The coefficient on the attic insulation variable measures the impact of insulation on energy use for "normal" weather conditions, conditional on other characteristics of the household and house. In equation (10'), we allow for a house specific intercept to capture time invariant characteristics of the house and residents not measured in the data. Below we will contrast this approach to an approach where we explicitly include temperature conditions as explanatory variables and use actual energy consumption rather than NAC. In that case equation (10') is replaced by

(10") $$\ln Q_{it} = \alpha_i + \beta_1 K_{it} + \beta_2 X_{it} + \beta_3 W_{it} + \beta_4 W_{it} K_{it} + \epsilon_{it}$$

where Q is actual energy consumption and W is a vector of heating and cooling degree days as well as squared heating and cooling degree days. We also interact temperature conditions with investment variables in this regression.

We begin by reporting results for the non-conservation variables from a regression of normalized energy consumption (Table 5).[11] Energy consumption goes up with family size and with household income. For the latter, the dummy variable for low income households (below $10,000) is excluded. Energy use tends to increase with income with the highest income group using 23% more energy on average than the lowest income group. Two measures of house size, area heated and number of windows, are strongly significant determinants of energy use. Houses with older furnaces tend to use more energy as do older houses. Houses built after 1979 on average use 24% less energy than houses built before 1940. The price measures are statistically significant; below, we report estimates of elasticities evaluated at mean energy prices. There is a downward trend in energy consumption in the data on the order of 1.8% per year. All in all, the non-conservation results are quite plausible.

Table 6 provides the first results for attic insulation. The first row reports results from the pooled data in which we assume a common intercept across all households. The first column result in that row is from the regression reported in Table 5. This coefficient estimate has the incorrect sign (insulation associated with increased energy consumption) and is statistically insignificant. Before considering issues of endogeneity and correlations of the investment variable with the error term, we consider the importance of sample selection. Column 2 excludes households for which there is a very large change in NAC between samples (greater than 50% in absolute value). Excluding these observations reduces the estimated coefficient somewhat but it is still positive and statistically insignificant. If we limit the analysis to single family homes, the estimated coefficient is now negative and nearly twice its standard error. According to this

[11] The regression also includes the attic insulation dummy. This regression does not include a house specific intercept.

estimate, attic insulation reduces energy use by 3.3% on average. That results should change so much for selection on single family units is not very surprising. There is likely to be considerable noise in the relationship between energy conservation and consumption for trailers, multi-family units, and other attached units[12]. If we also exclude large NAC change outliers (last column), the estimated coefficient increase slightly (-3.9%). Clearly, the single family restriction is the most significant one.

In the next row, we difference the data to remove individual specific intercepts from the data. If these intercepts (reflecting unobserved household and householder characteristics) are correlated with the insulation variable, then the coefficient estimate will be biased. There are two sorts of correlations to be concerned about. First, certain households may consume more energy (e.g. drafty houses) in ways not controlled for in our data (age of house, number of windows). These houses will likely have larger intercepts and are more likely to invest in attic insulation. This will induce a positive correlation between the intercept and attic insulation variable. Not controlling for individual specific intercepts then will bias the attic insulation variable in a positive direction. Second, certain householders may have a taste for energy conservation not captured by the demographic variables in the data set. These observations will likely have a lower intercept and a higher probability of attic insulation. This will induce a negative bias to our insulation coefficient if we do not control for individual specific intercepts.

When we do not restrict the data to single family houses, the control for individual specific intercepts changes the estimated coefficients quite sharply. They are now negative and quite large relative to their standard errors. Restricting ourselves to single family households, the coefficient

[12] The non-conservation coefficient estimates are not affected by the sample selection.

estimate falls substantially (from -.033 to -.059) and the t statistic exceeds 2. The results in row 2 suggest that controls for house specific intercepts helps mitigate problems of endogeneity in the attic insulation variable.

Next, we report results from a first difference regression in which all time varying variables (other than attic insulation) are excluded from the regression. In all cases, the coefficient estimate becomes larger (in absolute value) than the first difference regressions with controls. These results suggest the importance of controlling for as much house and family specific variation as possible in the data.

The final row provides two stage least squares estimates from the levels regression with instruments for the investment variable. We use two sets of instruments. The first set are education dummies indicating the level of education for the head of household. These instruments are motivated by the belief that education is likely to affect the decision to invest in conservation capital but is less likely to affect energy consumption itself (after controlling for household income). The second instrument set includes information about housing tenure and includes information on whether the family moved into the house prior to 1974 and whether they've moved in within the past 5 years. If householders moved into their house prior to 1974, they may have increased their attic insulation after the first oil shock in late 1973. Therefore, they may be less likely to put in additional insulation during our sample period. Also, one would expect householders to carry out large scale home improvement projects at the time (or soon after the time) when they move into a new house. None of these variables should be correlated with the error term in the energy consumption regression.

The last row of Table 6 reports the IV results. Before discussing the results, we point out that the first stage regression results support our conjectures about the investment decision. People who moved into their current house after 1974 are slightly more likely to have attic insulation though neither the effect nor its t-statistic are very large. However, householders who moved in less than 5 years ago are less likely to have attic insulation (coefficient estimate is -.07 with a t-statistic of 3.1). Finally, education variables suggest that more education is associated with having insulation in the attic though only the dummy variable for high school graduates is statistically significant (relative to high school dropouts).

Despite the conformance to expectations in the first stage regressions, the IV regressions perform quite poorly. None of the specifications have the expected sign nor are they statistically significant once outliers are excluded from the analysis. This may just reflect the fact that our instruments do not provide us with enough power to identify the effects of home improvement. In any case, the IV runs provide very little support for the view that there is an Energy Paradox.

We next present alternative measures of the impact of attic insulation on energy consumption. Rather than constructing a measure of normalized energy consumption, we use actual energy consumption and add heating and cooling degree day measures on the right hand side. We add HDD and CDD as well as HDD^2, CDD^2, and HDD and CDD interacted with the insulation variable. Column 1 reports results for the model in which we ignore the presence of correlated fixed effects while the next two columns incorporate fixed effects. The last column treats the investment decision as endogenous and instruments for it.

In both the levels and the first difference regression in which other variables are included, each insulation variable is imprecisely estimated by itself. However, the impact of insulation evaluated at mean levels of HDDs and CDDs is statistically significant with a mean response of

-7.2% for the levels regression and -5.8% for the first difference regression. The lack of statistical significance for the variables individually may simply indicate a problem of multicollinearity among the three variables, but it in no way affects their consistency or the precision of the estimated impact response. As in the NAC regressions, dropping other household and family specific information from the regression (third column) biases upward (in absolute value) the estimated impact of insulation on energy consumption. The IV regressions perform better in this specification with the mean energy consumption reduction reduced to 4.1%. It is imprecisely measured however with a t-statistic of 1.2. Both the first difference and the IV regressions generate coefficient estimates smaller (in absolute value) than the OLS levels regression. If this is a better specification than the NAC regression approach, then it appears that the correlation between energy use and a taste for conservation biases upward (in absolute value) the measured response of energy savings to the investment decision.

The closeness of the coefficient estimates in the first difference regressions when ln(NAC) is the dependent variable and when ln(Heat) is the dependent variable gives us some confidence in this coefficient estimate. The last step of the analysis in this paper is to convert this estimate into a percentage return. Before doing this, we consider other information contained in the regression findings, in particular estimates of price elasticities of demand.

VII. Estimates of Price Elasticity of Demand

We next turn to estimates of the price elasticity of demand generated from our micro data. Recall that we have three price variables. Price1 is a price for electricity for households whose main heating source is electricity. Price2 is a price for natural gas for households whose main heating source is gas while price3 is a price for electricity for those households. These prices are generated by dividing total expenditure for each fuel type by BTU consumption. Hence these are average prices and elasticity estimates will be biased to the extent that there is significant differences between the marginal and average price of energy. However, to the extent that households make consumption decisions based on the size of their energy bill rather than the marginal price, then these elasticity estimates will measure the correct responsiveness of consumers to changes in energy prices[13].

The first row of Table 8 gives the estimate of the price elasticity of electricity for households who heat with electricity (roughly 1/4 of the sample). The elasticity is quite high and exceeds 1 for the OLS regression where normalized energy consumption is the measure of energy consumption. Regardless of the method of estimation, this elasticity estimate is quite precisely estimated. For the normalized energy consumption regressions, the price elasticity of demand for electricity is between -.73 and -1.16 for households that heat with electricity and near zero for householders that heat with natural gas. A recent study on the demand for residential electricity (Branch (1993)) cites a number of studies in which the estimated price elasticity of demand ranges from -0.11 to -0.55. Our estimates are disaggregated by heating source and strictly speaking are not comparable to the estimated elasticities reported in Branch. The last row in Table 8 reports a weighted elasticity of demand for electricity where the estimates of Price1 and Price3 are weighted by the number of observations in each group. These estimates are more comparable to those reported in Branch and lie in the range of estimates he reports.

Finally, the second row reports price elasticities for natural gas. These elasticities are

---

[13] Wilder and Willenborg (1975) and Shin (1985) present evidence that consumers respond to average price rather than to marginal price.

uniformly lower than the elasticities for electric heating (price1). Herbert and Kreil (1989) report results from which a price elasticity of demand for natural gas can be computed which is equal to -0.36. Houthakker and Taylor (1970) report a short run price elasticity for natural gas of -.15. Our estimate is -0.47 for the OLS regression and -0.48 for the IV regression. The first difference regression is much smaller (-0.03) but imprecisely estimated.

The heat regressions tell a similar story. The electricity demand elasticity for households that heat with electricity are substantially larger than for natural gas heating households and the latter elasticity measure is typically statistically insignificant. However, for the OLS regression, it is positive and significant, an unexpected result that may suggest the advantage of the NAC approach over the heat regression approach. However, this positive elasticity disappears when the individual effect is eliminated. Across the board, the elasticity estimates for electricity using the actual energy consumption are greater (in absolute value) than when normalized energy consumption is used. On the other hand, the natural gas estimates are smaller. None of the differences are substantial across regressions and the weighted electricity elasticities are quite similar.

VIII. Returns to Conservation

We next turn to converting the energy savings estimates in the regressions to dollar savings to estimate a measure of the economic return to conservation. To do this, we must make assumptions about the future path of energy prices and obtain an estimate of the costs of attic insulation. Given assumptions about the future path of energy prices, we can compute the internal rate of return on the conservation investment. Let K be the cost of the investment and C be current energy cost savings ($C = \hat{\beta} pE$, where p is the price of energy and E is consumption in

BTUs or some other unit of energy use and $\hat{\beta}$ is the estimated regression coefficient for making the investment). Under the assumption that energy prices grow at a fixed rate γ and that energy consumption is constant, we can compute the internal rate of return (ρ) by solving the following equation:

$$\sum_{t=1}^{T} \frac{C(1+\gamma)^t}{(1+\rho)^t} - K = 0$$

If the returns continue forever ($T=\infty$) and $\rho > \gamma$, then the internal rate of return is given by

$$\rho = \frac{C(1+\gamma)}{K} + \gamma$$

The mean ratio of energy cost savings to costs of attic insulation (C/K) for households that make investments in our sample is 14.9%. If consumers assume that current energy prices will not change in the future ($\gamma = 0$) then the return on this investment is 14.9%. A return in this range is not particularly large and suggests that the engineering produced returns overestimate the returns to conservation. The return is higher (lower) if energy prices are expected to increase (fall) in the future. Table 9 presents some estimates of the return on attic insulation for different values of C/K and the growth rate of energy prices. Using the mean estimate of the current return (C/K) and assuming no growth in energy prices, then the discount rate implied by the data is bounded above by 14.9%. Table 9 also presents a range of estimates based on different assumptions in the growth rate of energy prices and the mean current return. The lowest estimated discount rate implied by the returns in Table 9 is less than 4% while the highest is less than 30%. This highest estimate begins to lend credence to the existence of the Energy Paradox; however this estimate

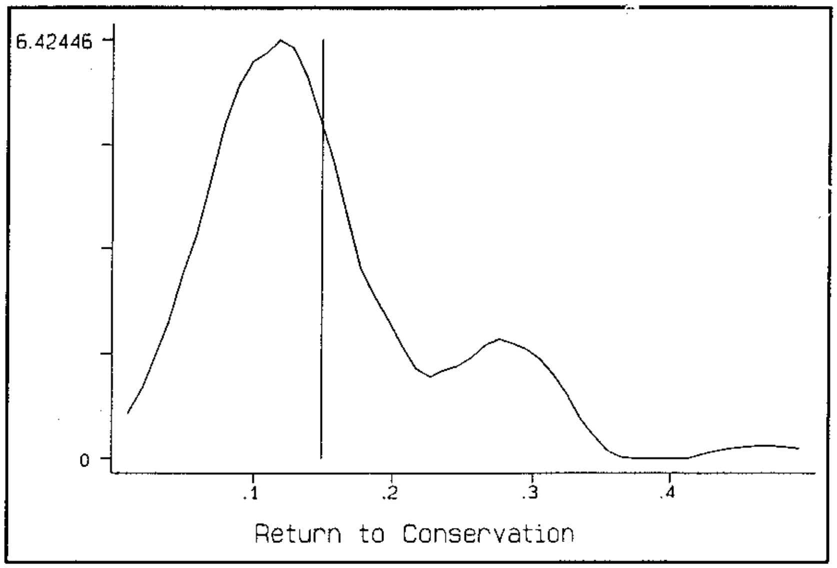

Figure 1

requires a fairly high current return (nearly 24% per year) plus a 3% real growth in energy prices. Note also that the mean return (14.9%) is from a highly skewed distribution of returns in our sample. Figure 1 above shows the distribution of returns in our sample among those households that actually invested in attic insulation. The vertical line is drawn at the mean return. The median return is 12.3% and over 75% of the households in the sample have a return less than 18%. Thus, while Table 9 suggests the possibility of the Energy Paradox using the mean return plus 1 standard deviation, the distribution of returns in our sample suggest that for the bulk of investors, the realized return is far below that required for the existence of an Energy Paradox.

A more complete comparison of the realized returns from this study and underlying discount rates must control for the risk of the investment. We can think of the investment as a project and calculate the project beta from which we can compute the risk adjusted discount rate. The return on this investment is based on energy savings (fuel oil, natural gas, and electricity).

The beta for crude petroleum and natural gas is on the order of 1.1 and for electric utilities of .5[14] Given the fraction of houses that heat with natural gas and with electricity in our data set, these values of beta would suggest a project beta of roughly 1 for our sample. A beta of 1 means that the project discount rate should match the market return, a number on the order of 7 to 10 percent.[15] If this is correct, then we nearly have shown that there is essentially no Energy Paradox for this investment.[16]

Before concluding that there is no Energy Paradox, we should last check whether the households that invest receive a greater return than do those that do not invest. If not, one might conclude that there is still the possibility of an Energy Paradox since the non-investors are passing up an investment that has a return higher (albeit not very much) than the appropriate discount rate. We can not measure the return for those households that do not invest. But we can look at attributes of the household that might correlate with the return. Means for these variables broken down between investors and non-investors are presented in Table 10. First we report statistics for the age of house. Older houses are expected to benefit more from insulation while newer houses are more likely to be well insulated upon construction. The test statistics bear this out. Houses for which investments are not made tend to be newer. Roughly 50% of the houses in which

---

[14] Brealey and Myers report an industry beta of 1.07 for crude petroleum and natural gas and .46 for electric utilities (Table 9-3, page 173).

[15] Kocherlakota (1996) reports a real return on the S&P500 of 7 percent. Feldstein, Poterba, and Dicks-Mireaux (1983) report a pretax real rate of return on capital for the postwar period of 12 percent. On an after-corporate tax basis, this would produce a return around 9 to 10 percent.

[16] This ignores the irreversible nature of the investment. Incorporating this would push the hurdle rate up even higher (see Dixit and Pindyck (1994) for a full treatment of this concept).

investments are made were built prior to 1950. Second, we look at energy consumption (using the normalized energy consumption measure). Houses with large amounts of energy consumption will benefit more from investment as a 6% energy reduction will mean greater reductions on their energy bill. The mean energy consumption for investors is over 10% higher than for non-investors. The test statistic for equality of means for energy consumption equals 2.00 suggesting that we can reject the hypothesis of equal means for investors and non-investors. Finally, we report statistics on the mean price of the main heating fuel for houses (per 1000 BTUs). Again, the return to the investment will be greater at higher energy prices. The results are less conclusive here. For households heating with electricity the mean price for investors is slightly higher but we cannot reject the hypothesis that the mean prices are the same across the two groups. For households heating with gas, the price of gas for non-investors is slightly higher but the difference is economically negligible.

We also considered other variables that one might believe are correlated with the decision to invest but which are not directly related to the return that households receive from conservation investment. Mean income, education levels, and age of main householder all show no statistically significant differences between investors and non-investors. Since these variables don't affect the return on the investment, we do not find these results surprising.

Summing up, the results from Table 10 lend support to the hypothesis that there is no Energy Paradox. Mean returns for investors is not significantly greater than one would expect from an CAP-M analysis of the energy conservation decision and differences between investors and non-investors suggest that non-investors would not have received higher returns had they invested than did those households that did invest.

IX. Conclusion

We combine monthly energy billing data with the annual Residential Energy Conservation Survey (RECS) in order to assess the energy savings from home improvement investments. We use the panel features of our data to control for individual and structure specific heterogeneity, and find that this step importantly affects our inferences. We find the following in this paper. First, unobserved and correlated heterogeneity among households is important to control for when estimating returns to conservation. Based on the heat regressions, it appears that the decision to invest in attic insulation is *ceteris paribus* negatively correlated with energy consumption. Failing to control for this correlation leads to an overestimate of the returns to conservation. Second, the data--- which may well be the most comprehensive yet applied to this question--- provide little evidence of an Energy Paradox. The mean rate of return in this analysis for attic insulation is 14.9% and the median is 12.3%. These rates put an upper bound on the implied discount rate for the energy investments analyzed in this paper and are consistent with plausible discount rates suggested by a CAPM analysis. In closing, we feel it is important to note that these results do not necessarily imply that subsidies for home improvement activity are bad policy. Even if consumers rationally account for all factors that directly affect their purchase of home improvement technologies, they might not account for the possible social costs of higher pollution associated with their energy consumption choices.

References

Arthur D. Little, Inc., "Measuring the Impact of Residential Conservation Programs: An Econometric Analysis of Utility Data," Final Report for RP-1587, Electric Power Research Institute, Vol. 8, 1984.

Berkovic, J. J. Hausman, and J. Rust, "Heating System and Appliance Choice," MIT Energy Laboratory Report No. MIT-EL 83-004WP, 1983.

Branch, E.R., "Short Run Income Elasticity of Demand for Residential Electricity Using Consumer Expenditure Survey Data," The Energy Journal, 14(1993): 111-121.

Brealey, R. and S. Myers, Principles of Corporate Finance, New York: McGraw-Hill, 1984.

Cole, H. and R. Fuller "Residential Energy Decision Making: An Overview with Emphasis on Individual Discount Rates and Responsiveness to Household Income and Prices," Hittman Associates Inc., 1980.

Dixit, A. and R. Pindyck, Investment Under Uncertainty, Princeton NJ: Princeton University Press, 1994.

Feldstein, M., J. Poterba, and L. Dicks-Mireaux, "State and Local Taxes and the Rate of Return on Non-Financial Corporate Capital," Journal of Public Economics 21(1983): 129-157.

Fels, M. "PRISM: An Introduction," Energy and Buildings 9(1986): 5-18.

Hassett, K. and Gilbert Metcalf, "Residential Energy Tax Credits and Home Improvement Behavior," Journal of Public Economics 57(1994): 201-217.

Hartman, R.S. and Michael Doane, "Taking the Con out of Conservation Program Evaluation," Resources and Energy, 9, 1987, 187-207.

Hausman, J. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables," Bell Journal of Economics, 10(1979): 33-54.

Herbert, J. and E. Kreil, "Specifying and Evaluating Aggregate Monthly Natural-Gas Demand by Households," Applied Economics, 21(1989): 1369-1381.

Hirst, E., "Actual Energy Savings After Retrofit: Electrically Heated Homes in the Pacific Northwest," Energy, 11(1986): 299-308.

______, "Energy and Economic Effects of Utility Financial Incentive Programs: The BPA Residential Weatherization Program," The Energy Journal, 8 (2), 1987, 97-110.

Hirst, E. and Richard Goeltz, "The Economics of Utility Residential Energy Conservation Programs: A Pacific Northwest Example," The Energy Journal, 5 (3), 1984, 159-169.

_____ ,"Estimating Energy Savings Due to Conservation Programs: The BPA Residential Weatherization Pilot Program," Energy Economics, January 1985, 20 -28.

Houthakker, H. and L. Taylor, Consumer Demand in the United States: Analyses and Projections, Cambridge: Harvard University Press, 1970.

Kocherlakota, N.R., "The Equity Premium: It's Still a Puzzle," Journal of Economic Literature, 34(1996): 42-71.

Sebold, F.D. and Eric Fox, "Realized Savings from Residential Conservation Activity," The Energy Journal 6 (2), 1985, 73-88.

Shin, J., "Perception of Price When Price Information is Costly: Evidence From Residential Electricity Demand," Review of Economics and Statistics, 67(1985): 591-598.

Train, K., "Discount Rates in Consumers' Energy Related Decisions: A Review of the Literature," Energy, 10 (12), 1985, 1243-1253.

Wilder, R. and J. Willenborg, "Residential Demand for Electricity: A Consumer Panel Approach," Southern Economic Journal, 42(1975): 212-217.

Table 1. Residential Energy Consumption Survey
Data Sample Construction

	Number of Households
Households in Data Set	6229
Primary Heat Gas or Electric	4544
Valid Monthly Billing Data	2524
Households Present In Both Years	**1136**
Sample Selection:	
(1) No NAC Outliers	**1076**
(2) Complete Household Data	**836**
(3) Single Family Units	**945**
(1), (2), and (3)	**765**

Valid monthly billing data are data for which at least 12 months of data are available and for which all billing periods are less than 70 days in length. Complete household data refers to households for which no missing data for regressors in equation (10) exist. Most missing data are for equipment age. NAC outliers are households which register a 50% or more increase or decrease in NAC between periods. See text for fuller discussion.

Table 2. Monthly Billing Data Coefficients

	1984	1987 Wave 1	1987 Wave 2	1990
HDD (β_1)	1.617 (1.003) .962 2.070	1.710 (1.171) .977 2.173	1.637 (1.092) .914 2.074	1.631 (1.043) .921 2.077
CDD (β_2)	1.189 (2.342) .311 1.587	.813 (1.767) .238 1.212	.879 (1.278) .255 1.265	1.083 (1.560) .340 1.507
Int (α)	2.635 (2.698) .931 3.880	2.914 (3.070) 1.175 4.035	2.757 (2.348) 1.346 3.869	2.656 (2.417) 1.085 3.894
N	622	622	514	514

Summary statistics on coefficient estimates from equation 8. The data for 1987 are broken down according to which wave the household is in. The top number in each group is the mean coefficient estimate. Its standard deviation is reported in parentheses underneath. The next two numbers are the 25th and 75th percentile values. Coefficients for HDD and CDD are change in million BTUs per change in 100 degree days. The coefficient for the intercept is monthly consumption in millions of BTUs.

Table 3. Normalized Annual Consumption (NAC)

	1984	1987 Wave 1	1987 Wave 2	1990
Mean	118.7	121.3	112.3	111.6
St. Dev.	60.8	60.4	61.1	60.1
25th Percentile	73.9	79.1	66.6	70.0
75th Percentile	153.6	154.5	143.7	143.9
N	622	622	514	514

This table reports the summary statistics on estimated normalized annual consumption (NAC) per household in the dataset. Units are millions of BTUs per year.

Table 4. Summary Statistics on Data

	N	Mean	SD	Minimum	Maximum
Log(NAC)	2272	4.60	0.62	1.38	6.01
Attic Insulation	2272	0.77	0.42	0.00	1.00
Family Size	2272	2.59	1.41	1.00	10.00
Non-White	2272	0.12	0.33	0.00	1.00
Hispanic	2272	0.04	0.21	0.00	1.00
Region					
New England	2272	0.03	0.16	0.00	1.00
Middle Atlantic	2272	0.10	0.30	0.00	1.00
E.North Central	2272	0.19	0.39	0.00	1.00
W.North Central	2272	0.19	0.39	0.00	1.00
S. Atlantic	2272	0.11	0.31	0.00	1.00
E.South Central	2272	0.10	0.29	0.00	1.00
W.South Central	2272	0.08	0.26	0.00	1.00
Mountain	2272	0.07	0.25	0.00	1.00
Pacific	2272	0.15	0.35	0.00	1.00
Household Income:					
less than $10,000	2272	0.16	0.37	0.00	1.00
$10,000-20,000	2272	0.22	0.41	0.00	1.00
$20,000-30,000	2272	0.18	0.38	0.00	1.00
$30,000-40,000	2272	0.17	0.38	0.00	1.00
$40,000-50,000	2272	0.10	0.30	0.00	1.00
$50,000-75,000	2272	0.13	0.33	0.00	1.00
more than $75,000	2272	0.05	0.23	0.00	1.00
Area Heated (ft^2)	2272	1,757	897	151	9,460
Number of Windows	2272	13.30	6.74	1.00	52.00
Age of Furnace:					
less than 2 years	1916	0.06	0.24	0.00	1.00
2 - 4 years	1916	0.13	0.33	0.00	1.00
5 - 9 years	1916	0.21	0.41	0.00	1.00
more than 10 years	1916	0.60	0.49	0.00	1.00

Table 4. Summary Statistics on Data (Continued)

	N	Mean	SD	Minimum	Maximum
Age of House					
before 1940	2272	0.23	0.42	0.00	1.00
1940s	2272	0.10	0.30	0.00	1.00
1950s	2272	0.17	0.38	0.00	1.00
1960s	2272	0.27	0.44	0.00	1.00
1970s	2272	0.21	0.40	0.00	1.00
after 1979	2272	0.03	0.16	0.00	1.00
Thermostat Setting	2204	69.96	3.84	50.00	84.00
Central or Room Air Conditioners	2272	0.70	0.46	0.00	1.00
Price1	545	20.18	5.08	6.95	34.84
Price2	1727	5.59	1.16	2.43	15.70
Price3	1727	23.74	6.06	5.11	57.99
Cooling Degree Days (x1000)	2272	1.27	0.86	0.03	5.29
Heating Degree Days (x1000)	2272	4.49	2.04	0.01	11.92

Table 5. Regression Results:
Non-Conservation Variables

	Coefficient Estimate	Standard Error	T statistic
Family Characteristics			
Family Size	0.064	0.006	11.193
Non-white	0.053	0.026	2.092
Hispanic	0.085	0.040	2.115
Household Income			
$10,000-20,000	0.031	0.026	1.198
$20,000-30,000	0.036	0.027	1.189
$30,000-40,000	0.091	0.028	3.258
$40,000-50,000	0.139	0.032	4.318
$50,000-75,000	0.133	0.031	4.263
more than $75,000	0.242	0.040	5.997
Region			
Middle Atlantic	-0.153	0.054	-2.847
E.North Central	-0.243	0.053	-4.574
W.North Central	-0.336	0.054	-6.213
S. Atlantic	-0.318	0.055	-5.794
E.South Central	-0.476	0.058	-8.149
W.South Central	-0.364	0.058	-6.326
Mountain	-0.356	0.057	-6.200
Pacific	-0.604	0.055	-11.021
Area Heated (x1000)	0.113	0.010	10.880
Windows	0.010	0.001	7.253
Furnace Age			
2 - 4 years	0.003	0.036	0.071
5 - 9 years	0.047	0.035	1.331
more than 10 years	0.059	0.032	1.848

Table 5. Regression Results:
Non-Conservation Variables (Continued)

	Coefficient Estimate	Standard Error	T statistic
Age of House			
1940s	-0.084	0.028	-2.952
1950s	-0.086	0.024	-3.619
1960s	-0.077	0.023	-3.289
1970s	-0.110	0.027	-4.050
after 1979	-0.240	0.054	-4.443
Air Conditioning Used	0.047	0.019	2.466
Thermostat Setting	0.011	0.002	5.105
Price1	-61.916	2.545	-24.324
Price2	-26.742	6.215	-4.303
Price3	-3.402	1.566	-2.173
Year	-0.018	0.004	-4.931
Intercept	5.990	0.348	17.204
N	1873		
Standard Error of Regression	0.314		
Adj. R^2	0.673		

Estimated coefficients from a regression of ln(NAC) on household characteristics and conservation capital (equation 10'). This regression is an OLS regression without estimation of individual fixed effects. The conservation coefficient estimate is reported in the next table.

Table 6. Regression Results:
Conservation Variables

	Full Set	No NAC Outliers	Single Family	Both Restrictions
Levels	0.023	0.013	-0.033	-0.039
	1.014	0.587	-1.363	-1.620
FD-1	-0.048	-0.044	-0.059	-0.045
	-1.664	-2.126	-2.038	-2.197
FD-2	-0.070	-0.064	-0.069	-0.062
	-2.182	-3.246	-2.395	-3.252
IV	0.522	0.420	0.079	0.039
	2.420	1.747	0.300	0.143

This table reports coefficient estimates for attic insulation dummy along with its t-statistic. Column 1 reports results for the full data set. Column 2 excludes observations for which the change in NAC exceeds 50% in absolute value between the two sample periods. Column 3 limits the sample to single family homes and column 4 imposes both restrictions. The first row regression does not incorporate individual specific intercepts. The next row differences the data to sweep out individual specific intercepts. The third row excludes time varying variables other than attic insulation. The last row instruments for the investment dummy variable.

Table 7. Heating Regressions

	Levels	FD-1	FD-2	IV
Insulation	-0.131 -0.900	-0.170 -0.880	-0.077 -2.616	1.589 2.177
Insulation* CDD	0.046 0.977	0.032 0.521	--	-0.457 -2.131
Insulation* HDD	-0.001 -0.016	0.015 0.567	--	-0.230 -2.354
Impact	-0.072 -2.544	-0.058 -1.833	-0.077 -2.616	-0.041 1.271

Regression results from equation (10") for the sample of single family households. The log of energy consumption (in BTUs) is the dependent variable. In addition to the household and family variables used in the previous regressions, we include heating degree days (HDD), cooling degree days (CDD), heating and cooling degree days squared and interactions between HDD and CDD with the insulation variable. The last row of the table gives the marginal impact of attic insulation evaluated at mean HDD and CDD for the sample.

Table 8. Price Elasticity Estimates

	NAC Regressions			Heat Regressions		
	OLS	FD-1	IV	OLS	FD-1	IV
Price1	-1.13**	-0.73**	-1.12**	-0.84**	-0.43**	-0.89**
Price2	-0.47**	-0.03	-0.47**	-0.70**	-0.13	-0.72**
Price3	0.001	-0.08	-0.01	0.21**	-0.04	0.19**
Weighted Elasticity	-0.28	-0.20	-0.30	-0.26	-0.12	-0.04

** - statistically significant at the 1% level.

This table reports price elasticities based on coefficient estimates from equations (10') and (10"). Row 1 (Price1) reports the price elasticity of demand for electricity for households who heat with electricity. The next row (Price2) reports the price elasticity of demand for natural gas for households who heat with natural gas. The third row (Price3) reports the price elasticity of demand for electricity for households who heat with natural gas. The last row reports an weighted price elasticity of demand for electricity based on responsiveness of households using electricity for space heating (Price1) and households using natural gas for space heating (Price3). See text for more details.

Table 9. Economic Returns to Conservation

	C/K		
γ	-1 SD	Mean	+1 SD
-3%	3.4	11.5	19.5
-1%	5.5	13.8	22.0
0%	6.6	14.9	23.2
1%	7.7	16.0	24.4
3%	9.8	18.3	26.9

Source: Authors' Calculations. See text for details.

Table 10. Mean Characteristics of
Investors and Non-Investors Related to Investment

	Non-Investors	Investors	Test Statistic for Equality of Means
Age of House:			
pre 1940	0.23 (0.01)	.34 (0.06)	1.99
1940-1949	0.11 (0.01)	0.14 (0.04)	0.83
1950-1959	0.20 (0.01)	0.21 (0.05)	0.37
1960-1969	0.28 (0.02)	0.26 (0.05)	-0.38
1970-1979	0.18 (0.01)	0.03 (0.02)	-6.15
1980 +	0.02 (0.005)	0.01 (0.01)	-0.25
Energy Consumption (NAC)	125.82 (1.98)	140.46 (7.05)	2.00
Price of Heat:			
Heat w/ Electricity	0.0196 (0.0003)	0.0202 (0.0020)	.31
Heat w/ Gas	0.0053 (0.00004)	0.0052 (0.0001)	-1.43

Mean characteristics for households in year prior to investment (if investment is made). The test statistic in the last column is approximately standard normally distributed under the hypothesis that the means for investors and non-investors are the same.