COLORADO SCHOOL OF MINES
EARTH • ENERGY • ENVIRONMENT

Division of Economics and Business
Working Paper Series

Additionality, Mistakes, and Energy Efficiency Investment

Ben Gilbert
Jacob LaRiviere
and Kevin Novan

Working Paper 2019-01
http://econbus-papers.mines.edu/working-papers/wp201901.pdf

Colorado School of Mines
Division of Economics and Business
1500 Illinois Street
Golden, CO 80401

June 2019

Colorado School of Mines
Division of Economics and Business
Working Paper No. **2019-01**
June 2019

Title:
Additionality, Mistakes, and Energy Efficiency Investment*

Author(s):
Ben Gilbert
Division of Economics and Business
Colorado School of Mines
Golden, CO 80401
bgilbert@mines.edu

Jacob LaRiviere
Microsoft, University of Tennessee and University of Washington
jlariv@microsoft.com

Kevin Novan
Department of Agricultural and Resource Economics
University of California, Davis
knovan@ucdavis.edu

ABSTRACT
Investment subsidies, like those for energy efficiency upgrades, often suffer from poor additionality. Using a novel theoretical model, we show when benefits from investments are uncertain, subsidy policies can also lead to inefficient investment decisions that don't pass the full information benefit cost test. We use a unique household-level dataset of energy efficiency audits and investments to test model predictions. We find that households are significantly more likely to schedule audits and installs when first moving into an existing home when, ironically, information uncertainty is greatest. We then build on the theoretical model to inform better policy design.

JEL **classifications: Q, D**
Keywords: Durable Goods, Learning, Electricity, Second Best

*We thank Klaas van 't Veld for the use of his fuzzy string matching algorithm which helped immensely in the construction of our dataset. We also thank Brittany Tarufelli for excellent research assistance

1 Introduction

The residential sector accounts for 20% of all energy consumed in the United States.[1] To reduce the fuel, capital, and environmental costs required to supply this energy, policymakers have long sought to increase the energy efficiency of the housing stock. While much of the focus has been on ensuring that newly constructed homes meet some minimum level of energy efficiency (e.g., state-level building codes and federal appliance standards), the effect of these minimum standards is limited by the fact that existing older homes, and the appliances in them, are incredibly durable.[2]

To reduce the amount of energy consumed in the stock of older homes, a substantial amount of money is poured into programs encouraging investment in energy efficiency upgrades. Much of this support comes from customer incentive programs offered by energy utilities.[3] Estimates from the utilities themselves suggest that the energy savings achieved by these programs are quite large. For example, electric utilities estimate that the $2.9 billion they spent on energy efficiency programs and incentives during 2017 alone will provide lifetime electricity savings of 137,298 GWh.[4]

However, a growing body of academic research casts doubt on these lofty savings estimates. For one, ex post empirical estimates find consistent evidence that the energy savings achieved by energy efficiency upgrades are smaller than ex ante predictions (Fowlie, Greenstone and Wolfram (2018), Burlig et al. (2017), Allcott and Greenstone (2017), and Metcalf and Hassett (1999)). Perhaps even more importantly, however, the energy savings achieved by energy efficiency incentives are often far less than predicted due to the fact

[1] See https://www.eia.gov/energyexplained/index.php?page=us_energy_use.

[2] Over half of the current U.S. housing stock was built prior to 1980 – before any state-level building codes established minimum energy efficiency requirements. For information on the age distribution of the housing stock, see https://www.census.gov/programs-surveys/ahs.html.

[3] These programs provide both financial incentives – e.g., rebates and subsidies for energy efficiency investments – as well as informational interventions – e.g., home energy audits. From 2013 through 2017, U.S. electric utilities alone spent $7.5 billion on residential energy efficiency programs. Similar types of financial support also come from the federal government (e.g., tax credits for home improvements; the Weatherization Assistance Program). From 2013 through 2017, the federal government spent $2.2 billion on tax credits for homeowners making energy efficiency improvements to existing homes and another $1 billion to upgrade low-income homes through the Weatherization Assistance Program. For information on federal tax expenditures, see https://home.treasury.gov/policy-issues/tax-policy/tax-expenditures. Utility expenditures on energy efficiency are reported in the EIA's Electric Power Industry Report (EIA-861), https://www.eia.gov/electricity/data/eia861/.

[4] For a comparison, total U.S. residential electricity consumption during 2018 was 1.46 million GWh.

that they suffer from poor additionality (Joskow and Marron (1992), Boomhower and Davis (2014)). That is, many of the subsidized energy efficiency upgrades would have occurred without any direct financial support.[5]

In this paper, we highlight that poor additionality of energy efficiency subsidies is an incomplete representation of the subsidies' inefficiency. Introducing a theoretical model of a household's decision to invest in energy efficiency upgrades, we demonstrate that an additional inefficiency can arise due to the fact that there is uncertainty surrounding the benefits achieved by performing the upgrades. Specifically, many installing households don't know how much money will be saved or how much comfort will be increased at time of upgrade. In this setting, a subset of households will make irreversible investment "mistakes" – they will invest in energy efficiency upgrades that would prove to be inefficient under full information. Using our theoretical model, we demonstrate that, in the presence of a program subsidizing energy efficiency, not only will a large share of participating households be inframarginal (i.e. "non-additional"), many of the "additional" participants will not be economically efficient participants.

To explore how investments in energy efficiency are affected by financial incentives, we begin by modeling a household's decision to participate in an energy efficiency program when the household faces uncertainty about its average energy bill. Consistent with typical energy efficiency programs offered by energy utilities, we model this decision as a two-step process. In the first stage, a household must decide whether to receive a potentially subsidized in-home energy audit (IHEA). Through an audit, homeowners receive expert advice regarding what types of energy efficiency upgrades could be performed on their home as well as estimates of the potential energy savings. Importantly, while the expert advice will reduce the uncertainty surrounding the magnitude of the potential benefits, households will still not know the true returns. In the second stage, homeowners must then decide which, if any, of the subsidized energy efficiency improvements to make.

While a number of studies explore households' decisions to participate in similar energy efficiency programs (e.g., Holladay et al. (2019), Allcott and Greenstone (2017), and

[5] Low levels of additionality have also be highlighted in a variety of other settings where subsidies are provided to encourage investment in energy saving technologies – e.g., subsidies for residential solar PV (Hughes and Podolefsky (2015)) as well as subsidies for hybrid vehicles (Chandra, Gulati and Kandlikar (2010)).

Palmer and Walls (2015)), our analysis differs in an important dimension. Rather than focusing solely on whether or not a household participates in an energy efficiency program, we seek to understand the timing of households' participation decision. That is, when do households receive audits and make energy efficiency upgrades? The timing of participation is potentially important given that households may have very different information based on how long they have lived in their homes. When new owners first move into a home, they will likely have a great deal of uncertainty surrounding how much future energy they will consume as well as what the resulting comfort level will be (i.e. will their home be drafty in the winter or too hot in the summer). Consequently, new homeowners will have the least certainty surrounding the benefits that would be achieved by making energy efficiency upgrades. In contrast, homeowners that have resided in their homes for a great deal of time will have a far more precise understanding of the benefits energy efficiency upgrades would provide.

To explore how uncertainty affects the homeowners' decisions to perform energy efficiency upgrades, we model bayesian households that form priors about a home's energy usage based on the observable characteristics of the home (e.g., the age and visible condition of the home) when they first move in. The households then update their priors with each new energy bill. In this context, our model identifies four different types of households: (1) households that correctly make energy efficiency upgrades, (2) households that correctly choose not to make upgrades, (3) households that delay making upgrades that would be economically efficient, and (4) households that make upgrades which are inefficient and don't pass a full information benefit-cost test. Our model suggest that the likelihood of making a mistake (type 3 or 4) is the greatest when uncertainty is the largest – right when households move into a home. Paradoxically, this is also precisely when our model suggests households are most likely to participate in energy audit and subsidy programs.

To explore whether the predictions from our model are borne out in practice, we examine the participation decisions of households residing in a medium-sized MSA which subsidized energy audits as well as subsequent energy efficiency upgrades. To perform the empirical analysis, we combine three unique datasets: one with monthly, household-level electricity and gas usage, one with audit and install information, and one with home characteristics. Using the data, we examine how the likelihood of participating in the audit program and making subsequent upgrades varies with characteristics of the premise (e.g., the age of the home) as well as with the time the homeowners have spent living in the home. Consistent with the prediction from our theoretical model, we first document that there is a dramatic increase in the likelihood of receiving an audit immediately after a household moves into a home. Moreover, consistent with new homeowners using observable premise characteristics to inform their beliefs about the potential benefits of energy efficiency upgrades, the likelihood of an audit immediately after moving in is higher in older homes. In line with the predictions from the theoretical model, the likelihood of an audit occurring falls precipitously with each additional month a household resides in their home. Moreover, the likelihood of an audit occurring is no longer a function of observables like home age once the homeowners have resided in the premise for several months. This final finding is consistent with the homeowners using the information they receive from actually living in the home (e.g., a series of energy bills; the temperature and draftiness of the home) – as opposed to simple observable premise characteristics – to inform their beliefs about the potential benefits of making energy efficiency upgrades.

With the empirical evidence supporting our theoretical model, we return to the bayesian model to explore how households' decisions are affected by policies subsidizing energy efficiency upgrades versus audits. Focusing first on the impact of subsidies for performing upgrades, our model demonstrates that, while many of the participating households are "infra-marginal", increasing the size of the subsidies will indeed cause more households to invest in upgrades. To understand the social returns provided by these new participants, we separate the marginal households into three groups. First, there is a set of households that are truly additional and efficient participants. The increase in the subsidy incentivizes these households to invest in upgrades that provide positive social returns. Second, there is a set of households whose investments are "pulled forward" in time.[6] While these households would have eventually chosen to make investments in energy efficiency upgrades in the future, the increased subsidy incentivizes them to perform the upgrades sooner. While

[6]Evidence of agents actions being pulled-forward in time has also been demonstrated in the case of subsidies for vehicle scrapping (Mian and Sufi (2012)) as well as for solar PV adoption (Hughes and Podolefsky (2015)).

pulling these investments forward provides social benefits, the returns are far smaller than the gains provided by the first group of truly additional households. Third, there is a set of marginal households that are incentivized to invest in energy efficiency upgrades that are not socially optimal. That is, the subsidy simply encourages more investment mistakes.

Turning our attention to the impact of subsidizing energy audits, we find an unfortunately similar result. By reducing the private cost of audits, additional households will certainly be encouraged to pay for audits. However, among the recent movers that do not yet know what their typical energy consumption levels will be, performing an audit will not be able to remove all of the uncertainty surrounding the benefits of making energy efficiency improvements. While we find that the increase in the audit-uptake does lead to more investment in energy efficiency improvements, many of the additional upgrades are again merely pulled forward in time while others are economically inefficient upgrades – i.e. the costs again exceed the true benefits.

The findings from our theoretical analysis first suggest that the social benefits provided by subsidizing energy efficiency upgrades may be smaller than previously thought. Existing studies demonstrate that only a fraction of subsidized energy efficiency upgrades represent additional investments (e.g., less than 50% in the setting explored by Boomhower and Davis (2014)). Our analysis highlights that many of these additional investments are not truly additional, but rather simply pulled forward in time. Even more troubling, given the existence of uncertainty in the decision making process, we highlight that many of the additional participants may be making inefficient investments – i.e. upgrades with costs larger than the stream of social benefits they will provide.

More generally, our analysis contributes to a growing literature examining how the provision of information affects consumers' investments in energy-related durables.[7] When purchasing an energy-consuming durable good (e.g., a vehicle, an appliance, or even a house), consumers face a trade-off between the upfront purchase price and the future stream of energy payments required to operate the durable good. Acknowledging that

[7] A related literature (e.g., Hausman et al. (1979) and Dubin and McFadden (1984)) explores whether consumers behave myopically when purchasing an energy-consuming durable good. That is, do consumers undervalue the future operating costs relative to the upfront cost? If the answer is yes, then this could be driven in part by downwardly biased beliefs about the savings provided by energy efficiency. Recent evidence focusing on vehicle purchases finds little (e.g., Busse, Knittel and Zettelmeyer (2013)) to no (Allcott and Wozny (2014)) systematic undervaluation of fuel efficiency.

inefficient investment decisions will occur without complete information regarding these future operating costs, policymakers have long focused on mandating the provision of information.[8] Recent work demonstrates that the provision of this information can meaningfully alter consumers' appliance purchase decisions (Houde (2018), Allcott and Taubinsky (2015), and Newell and Siikamäki (2014)). Moreover, providing more precise signals about savings can result in more efficient investment decisions (Davis and Metcalf (2016)). Our paper is somewhat different in that we investigate learning about the state of the world through multiple information draws rather than more precise signals.[9] This distinction is vital in our context since a new homeowner's information changes over time with each new bill that makes the timing of their investment decision important. The model permits us to highlight the policy importance of this distinction.

Our analysis suggests that, in the case of households investing in upgrades to their homes, meaningful welfare gains could be achieved if the uncertainty surrounding the benefits from improving a home's energy efficiency could be eliminated. In particular, a key contribution of this study is highlighting that this information would be the most valuable when provided to households at the time they move into their homes – when the likelihood of performing inefficient energy efficiency upgrades peaks.[10] At first glance, this mover margin appears to have a great deal of scope. Roughly 7% of existing U.S. homes are sold each year, suggesting that nearly half of the existing homes will turnover within a 10 year time period.[11] However, our results suggest the need for a great deal of caution for targeting home sales. Households will have the most uncertainty surrounding the benefits

[8]For example, the U.S. Environmental Protection Agency requires that new cars and trucks for sale have fuel economy "window stickers" prominently displayed. Similarly, the Federal Trade Commission requires manufacturers of major household appliances (e.g., refrigerators, water heaters, etc.) to display EnergyGuide labels.

[9]The distinction is important in both the investment under uncertainty literature (e.g.,Dixit, Dixit and Pindyck (1994)) and stochastic control problems with uncertainty. When learning the value of a parameter in a stochastic system, the task is easier when stochasticity decreases (the signal is more precise as in Davis and Metcalf (2016)). We investigate a scenario where the decision maker faces the same level of precision but can choose to wait to observe multiple signals before exercising an investment opportunity.

[10]The idea of targeting energy efficiency programs is not a new concept. For example, the federal government's energy efficiency mortgage program targets homeowners seeking to perform upgrades at the point of sale. For information on the energy efficiency financing options available, see `https://www.energy.gov/energysaver/incentives-and-financing-energy-efficient-homes/financing-energy-efficient-homes`.

[11]The St. Louis Federal Reserve reports the annual number of owner occupied housing units sold. During each of 2016, 2017, and 2018, 7% of the existing stock was sold.

from making energy efficiency upgrades immediately after moving in. If a traditional audit does not remove this uncertainty, then targeting energy efficiency programs at movers may simply induce a deceptively large uptake effect, which creates the appearance of an effective program. However, if many of the participants are simply pulled-forward or encouraged to make inefficient investments, much of the program spending will be largely wasteful. These insights suggests that policy makers should focus resources instead on alternative audit designs that work with households to study how their behavior interacts with their home characteristics to better estimate and predict their own usage.

In section two we introduce the theoretical model used to examine household-level investment decisions in the face of uncertainty. In section three and four we describe the empirical setting and data used to examine the testable implications from the theory. Sections five and six discuss the insights and conclusions our theoretical model provides with regards to subsidizing energy efficiency upgrades and home energy audits.

2 Theoretical Model

2.1 Model Framework

Our model focuses on two sources of uncertainty in the decision to receive an in home energy audit and install energy efficiency upgrades. First, households are uncertain about their home's true mean energy use μ, but can learn it over time by observing energy bills. Second, households are uncertain about the energy savings from making an install. The only way to learn about install savings is by performing an audit. Expected benefits of an install are the product of mean energy bills and savings from an install.

Households are endowed with an exogenous level of wealth w in each period which must be allocated between exogenous energy usage (e_t) and consumption of a numeraire good (c_t). We assume risk neutral utility with a constant, exogenous price of energy (p) such that: $u(c_t) = w - pe_t$. Energy consumption in any time period t is an *i.i.d.* random variable with a distribution $f(e) \sim N(\mu, \tau)$ and associated CDF $F(e)$. We define this distribution in terms of precision τ rather than variance to simplify notation for the Bayesian updating process below (i.e. $\tau = \frac{1}{\sigma^2}$).

To make the subsequent analysis as transparent as possible, we have imposed several restrictive modeling assumptions. However, we highlight that these assumptions are innocuous in our setting. In particular, we discuss below that linear utility does not impact the sign of our comparative statics but facilitates closed form solutions.[12] In addition, receiving no direct utility from energy means that consumption decisions are not strategic with respect to updating and learning. We view this as likely in our scenario – households are not likely to strategically use energy upon moving into a new home in order to learn about their mean energy use in that home.

Time in our model is defined relative to the date a household moves into a home. In the initial time period, a household moves into a home and forms priors over average energy usage associated with that home. We assume priors over mean energy use μ are a function of a set of characteristics (e.g., the age of home, size of the home, etc.) households observe when they move into a home. We represent these home characteristics with a scalar θ – which then feeds into the household's prior estimate for mean energy usage, $m(\theta)$. Home characteristics vary across dwelling types such that initial priors also vary in the population according to an atomless distribution $m(\theta) \sim M(m(\theta))$.

For an individual household moving into a specific home type θ, we assume their priors are unbiased estimates of their true mean. More precisely, we assume that any individual household's true mean is distributed according to $\mu \sim N(m(\theta), r) = h(\mu)$ with associated CDF $H(\mu)$ and prior precision r. If prior precision r is higher, then observable characteristics of the home tightly predict subsequent energy usage, and vice-versa. Both $\mu(\theta)$ and $m(\theta)$ are increasing in θ (e.g., average energy use is greater in older or larger homes, ceteris paribus). To summarize, a household moving into a particular dwelling type θ forms a prior estimate $m(\theta)$ about mean energy use, which is an unbiased estimate of their true mean $\mu(\theta)$, around which energy use in each period e_t is distributed.

Households update their prior with each new observation of an energy bill as Bayesians. After t periods, assuming normal energy usage and normally distributed priors gives the following closed form posterior belief about mean energy:

$$m_t(\theta)|_{e_1,\ldots,e_t} \sim N\left(\frac{\tau m(\theta) + tr\bar{e}}{\tau + tr}, \tau + tr\right) \tag{1}$$

[12]For a summary fo the literature exploring how uncertainty and risk aversion may affect the decision to make energy efficiency investments, see Gillingham and Palmer (2014).

where $\overline{e}$ is the average observed energy use over the t periods. The household's belief about their mean use becomes more heavily weighted towards their observed historical sample average with less weight on their initial prior. In addition, the precision increases over time so that a homeowner who has only recently moved in to their home will have more uncertainty about their mean energy use than a homeowner who has been there for a number of months or years. We denote the household's posterior distribution for the mean at time t as $H_t(\mu)$.

Energy efficiency installs are available at a fixed cost F, which reduce energy bills by $(1-\alpha)\%$ for some $\alpha \in (0,1)$. After making an install, the household has energy bills of $\alpha p e_t$ and utility $U = w - \alpha p e_t$. We assume α is unknown unless a household has an audit. Consistent with most audit programs, households can schedule an audit at any time for a fixed cost of A. We assume that households have beliefs about α according to the distribution $G(\alpha)$. We assume $G(\alpha)$ is shared across all households.

2.1.1 Model Extensions

The assumption of a common $G(\alpha)$ is somewhat restrictive. An extension would be to allow some homes to have a higher α than others. However, different mean install savings $(1-\alpha)$ might be correlated with observable characteristics of a home and thus expected mean energy use $(m(\theta))$. For example, older homes have a higher mean energy use and higher percentage savings from an install. We discuss briefly below that allowing correlation between $(1-\alpha)$ and $m(\theta)$ would not impact the qualitative findings so long as the correlation is positive.

We've assumed that fixed costs for making an install are fixed over time. Previous studies present evidence that a sizable portion of the fixed causes of participating in energy efficiency programs are hidden, non-pecuniary costs (e.g. Fowlie, Greenstone and Wolfram (2015) and Allcott and Greenstone (2017)). In practice, the non-pecuniary fixed cost of making an energy efficiency installation may increase with the time spent in a home after the move-in date. For example, some energy efficiency upgrades, such as installing new attic insulation, are easier to perform before all of one's belongings are moved in. We also discuss the implications of allowing the fixed cost of making an energy efficiency installation to vary over time spent in a home F_t below.

2.1.2 Model Timing

Figure 1 shows the timing of the model. Initially, we assume a household must get an audit before making an install. This reflects the structure of the audit program we study empirically in the subsequent sections. In section 5, we explore the impacts of allowing households to make subsidized installs without first receiving an audit.

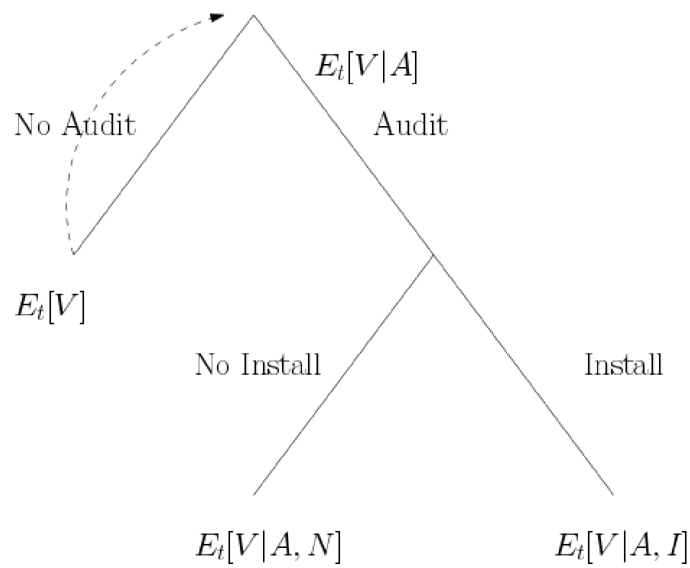

Figure 1: No audit, audit and install decision tree for each time period.

In each time step, a household can get an audit (A) or not. If a household doesn't get an audit in the current period, they can get an audit in the next period, represented by the curved dotted arrow. Conditional on getting an audit, the household can either make an install (I) or not make an install (N). For simplicity we assume households make a "yes" or "no" install decision in the same time step as when they get an audit. A household makes an install if the expected net present value (NPV) of doing so is positive. A household gets an audit if the expected value of doing so, incorporating the probability that it may lead to an install, is larger than the expected value of not doing so. Thus, expectations over savings from making an install matter for the household's audit decision. Finally, we assume that households passively update their information; in other words, households do not consider how future updating may alter today's audit and install decisions.[13]

[13]An alternative would be to fully model audits and installs as an exercise in option value. Given our

2.2 Analysis

We proceed by comparing expected value functions at each node in the decision tree shown in Figure 1. We start with the decision to make an install post-audit. We then backward induct to the audit decision.

Once the household has paid A for an audit and learned the install savings $(1-\alpha)\%$, they compare their expected value following an install against their expected value of doing nothing. Formally, the expected value of each option is:

$$E_t[V|A,I] = -F + \sum_{s=0}^{\infty} \delta^s \int w - \alpha p e_s \, dF(e(\mu)) \, dH_t(\mu) = -F + \frac{w - \alpha p m_t(\theta)}{1-\delta} \quad (2)$$

$$E_t[V|A,N] = \sum_{s=0}^{\infty} \delta^s \int w - p e_s \, dF(e(\mu)) \, dH_t(\mu) = \frac{w - p m_t(\theta)}{1-\delta} \quad (3)$$

Recall, $F(e(\mu))$ is the CDF governing energy bills and $H(\mu)$ is the CDF of the household's priors about the mean of their energy bills μ. The term δ is the time discount factor. $E_t[V|A,I]$ is the household's expected value function with an install, conditional on having had an audit and knowing the value of α. The household pays the install cost F up front but saves $(1-\alpha)$ in every period. The subscript t indexes the information available at time t about the distribution of energy bills. $E_t[V|A,N]$ is the household's expected value function if they do not make an install, conditional on having had an audit.

The expression for $E_t[V|A,I]$ reveals there is a critical value of α – which we define as $\bar{\alpha}(m_t)$ – such that the household is indifferent between making an install and not, given their beliefs.

$$E_t[V|A,I] = E_t[V|A,N] \Longrightarrow \bar{\alpha}(m_t) = 1 - \left(\frac{F}{p}\right)\left(\frac{1-\delta}{m_t}\right) \quad (4)$$

Thus, households make an install if $\alpha \leq \bar{\alpha}(m_t)$.

Equivalently, there is a critical belief threshold about mean energy use that a household would need to exceed in order to justify making an install for any given combination of

assumption of unbiased priors over mean energy use and linear utility, the gains of modeling fully forward looking households are low.

install cost (F) and install savings $(1-\alpha)$. Once an audit is performed and α is known, rearranging the expression in (4) provides the *ex post usage threshold*:

$$\bar{m} = \left(\frac{F}{p}\right)\left(\frac{1-\delta}{1-\alpha}\right),$$

such that it is rational to make an install if $m_t(\theta) \geq \bar{m}$.[14] In sum, these modeling assumptions have the intuitive implication that households are more likely to make installs if the share of energy saved by making them, $(1-\alpha)$, is larger for a given average bill, or their expected average bills are larger for a given percentage savings.

Stepping backward to the audit decision, there are two possible outcomes following an audit: make an install or not (realized α is below or above the critical value). Combining equations (2) and (3), the expected value function with an audit is:

$$E[V_t(audit)] = -A + \underbrace{\Big(1-G(\bar{\alpha})\Big)\int_e \frac{w-pe}{1-\delta}dF(e(\mu))dH_t(\mu)}_{\text{Probability-weighted value with no installation}} + \underbrace{\int_0^{\bar{\alpha}}\int_e \left[\frac{w-\alpha pe}{1-\delta} - F\right]dF(e(\mu))dH_t(\mu)dG(\alpha)}_{\text{Expected value with installation over range of savings}}. \quad (5)$$

Equation (5) shows that the expected value of getting an audit includes the value if no install is made, weighted by the probability that α is above the critical level $(1-G(\bar{\alpha}))$, plus the value with install savings integrated over the support of α's in which an install is made. For convenience we define the mean α conditional that an install would be made as $\widehat{\alpha} = E(\alpha|\alpha \leq \bar{\alpha}) = \int_0^{\bar{\alpha}} \frac{\alpha}{G(\bar{\alpha})}dG(\alpha)$. Then (5) can be simplified to

$$E[V_t(audit)] = \frac{w}{1-\delta} - A - F \cdot G(\bar{\alpha}) - \Bigg(1 - G(\bar{\alpha})(1-\widehat{\alpha})\Bigg)\frac{pm_t}{1-\delta}. \quad (6)$$

[14] If an audit was not required in order to make an install, households may make install decisions with uncertainty over α. In this case, we could define an *ex ante usage threshold* at which they would expect to make an install based on the distribution of α, $G(\alpha)$. This *ex ante usage threshold* is

$$\widehat{\bar{m}} = E(\bar{m}) = \int_0^1 \frac{F}{p}\frac{1-\delta}{1-\alpha}dG(\alpha).$$

This is the present value of income, less the fixed audit cost, less the expected fixed install cost, less the expected energy bills. These expected energy bills incorporate the probability of making an install and the expected savings if an install is made. Note that if the probability of making an install, $G(\bar{\alpha})$, is equal to one, then the expected energy bills are just the post-savings energy bills $\frac{\widehat{\alpha} p m_t}{1-\delta}$, where in that case $\widehat{\alpha}$ would be equal to the mean of the α distribution.

At the beginning of each time period, the household can choose to get an audit and receive $E[V_t(audit)]$ in equation (6), or delay one period and face the same choice again. The household takes the action with the highest expected value based on their priors of both α and μ. We therefore need to characterize the household's expected value if they delay one period and compare it to equation (6).

If the household delays the decision one period, they get the expected bill based on their current prior, plus the discounted ex ante expected value function:

$$E[V_t(delay)] = w - pm_t + \delta E[V_{t+1}],$$

where $E[V_{t+1}]$ is the household's ex ante expected value function, looking forward to node 1 in the subsequent period before the $t+1$ prior has been updated or the $t+1$ audit decision has been made. The node 1 decision depends on the distribution of the uncertain mean μ, $H_{t+1}(\mu)$.

Households will choose a costly audit if they think μ is large enough and α small enough to make an ex post install likely. This implies an ex ante critical belief about mean energy use at which the household is just indifferent between choosing the audit this period versus delaying the audit decision. We define this critical belief implicitly below, and denote it $\widetilde{m}$. To conserve notation, define $\widehat{\mu}_H = E(\mu|\mu \geq \widetilde{m}) = \int_{\widetilde{m}}^{\infty} \frac{\mu}{1-H(\widetilde{m})} dH(\mu)$ as the expected use conditional on use being above the cutoff to induce an audit. Similarly define $\widehat{\mu}_L = E(\mu|\mu < \widetilde{m}) = \int_{-\infty}^{\widetilde{m}} \frac{\mu}{H(\widetilde{m})} dH(\mu)$ as the expected use conditional on use being below the cutoff.

Using this notation, we show in appendix A.1 that the ex ante value function for $t+1$

is given by

$$E[V_{t+1}] = \underbrace{\frac{w}{1-\delta}}_{\text{present value of wealth}} - \underbrace{A\cdot\Big(1-H(\widetilde{m})\Big)}_{\text{ex ante expected audit cost}} - \underbrace{F\cdot G(\bar{\alpha})\Big(1-H(\widetilde{m})\Big)}_{\text{ex ante expected install cost}}$$

$$- \underbrace{\frac{p}{1-\delta}\left(1-G(\bar{\alpha})\cdot(1-\widehat{\alpha})\right)(1-H(\widetilde{m}))\widehat{\mu}_H}_{\text{expected bills if audit, given uncertainty over install}} - \underbrace{\frac{p}{1-\delta}H(\widetilde{m})\widehat{\mu}_L}_{\text{expected bills if don't audit}} \quad . \tag{7}$$

This is the present value of wealth net of what may happen with audits, installs, and energy bills after updating occurs. The second term (ex ante expected audit cost) accounts for the possibility that updated posterior beliefs about energy usage may induce an audit. The third term (ex ante expected install cost) accounts for the possibility of making an install if the updated belief does induce an audit. The last two terms provide the weighted average present value of bills, given that the updated posterior belief may be above or below the audit threshold.

Our empirical predictions about audit decisions therefore depend on the sign of the following expression:

$$E[V_t(audit)] - E[V_t(delay)] \gtreqless 0 \Longrightarrow$$

$$\frac{\delta}{1-\delta}\Bigg[\underbrace{\Big(1-G(\bar{\alpha})\cdot(1-\widehat{\alpha})\Big)(1-H(\widetilde{m}))p\widehat{\mu}_H + H(\widetilde{m})p\widehat{\mu}_L}_{\text{Expected future bills if delay, potentially audit and install later}} - \underbrace{\Big(1-G(\bar{\alpha})\cdot(1-\widehat{\alpha})\Big)pm_t}_{\text{Expected future bills if audit now}}\Bigg]$$

$$+ \underbrace{G(\bar{\alpha})(1-\widehat{\alpha})\cdot pm_t}_{\text{Expected savings this period}} - \underbrace{\Big(A+F\cdot G(\bar{\alpha})\Big)\cdot\left(1-\delta\Big(1-H(\widetilde{m})\Big)\right)}_{\text{Expected change in fixed costs}} \gtreqless 0. \tag{8}$$

Noting that $\bar{\alpha}$ and $\widehat{\alpha}$ are functions of the current belief m_t, equation (8) implicitly defines the critical belief $\widetilde{m}$. For beliefs above $\widetilde{m}$, the household audits in the current period. For beliefs below $\widetilde{m}$, the household delays and waits for more information.

Equation (8) also shows the household must account for the possibility of auditing and installing in the future when evaluating the expected costs and benefits of auditing now versus delaying one period. If the household delays, they will update their prior with their new energy bill. By the assumption of unbiased priors, their best expectation of tomorrow's updated prior is today's prior, m_t. However, the household assigns some probability $(1 - H(\widetilde{m}))$ to a state of the world in which their updated future prior will induce an audit in the future, which will potentially result in an energy efficiency install. Likewise, the household attaches probability $H(\widetilde{m})$ to a state of the world in which their updated future prior does not induce an audit. The household is more likely to audit in the present if expected current savings are large, if expected future bills are larger if delay rather than audit is chosen, and if the expected savings in fixed costs are low because the probability of eventually auditing in the future is high.

Thus, in the model, a household accounts for the various outcomes which can occur in addition to the possibility of making future audits. A household compares expected future bills to the expected bills if the household audits now and applies expected savings given their current prior m_t in all future periods. The household then compares expected future gains or losses to the expected savings in the current period from receiving an audit (and potential install). In addition, the household accounts for the fixed costs from auditing now with certainty and making an energy efficiency installation with probability $G(\bar{\alpha})$ versus potentially auditing and installing in the next period with probability $(1 - H(\widetilde{m}))$.

Equation (8) has several empirically testable implications, especially for households with more uncertainty in the beliefs over energy bills like recent movers. First, we show that audits are more likely soon after moving into a home; the updating and learning process leads to a declining share of households requesting audits as time passes from the move-in date.

Proposition 1 *The share of households receiving an audit in period t, rather than delaying one period, is declining in t.*

Two intuitive features drive Proposition 1 (a proof is in appendix A.2). First, recent movers have wide priors over their mean energy bills. As beliefs become increasingly precise over time, if a household has not audited yet because their belief is below the critical value then it is increasingly likely that their true mean is below the critical value. Second, there is attrition in unaudited households; at $t = 0$ all households with initial priors above the cutoff request an audit, and the share of households that subsequently updates their priors above the cutoff declines over time. As time passes, beliefs m_t converge to the true mean μ for each household, and an increasing share of households whose true $\mu > \widetilde{m}$ will have already audited earlier. Empirically, Proposition 1 indicates that we should observe the percentage of households receiving audits to be highest among recent movers and declining in time since the move.

Note that there are other possible models which give the prediction of recent movers auditing at high rates. One is that the cost of making an install decreases over time after moving. For example, it is easier to work on a house before being completely moved in. However, our other predictions below cannot be explained by time variable installation costs.

We are also interested in how households use observable information (i.e. home characteristics and recent energy bills) to make inferences about mean energy use over time, which ultimately affects audit and installation decisions. Observable building characteristics (θ) such as home age and size as well as large recent bill shocks e_t can influence the decision to audit rather then delay. An implication of Proposition 1 is that for every unaudited household as of time t, the probability that they will receive an energy bill large enough to update their posterior above the critical belief declines as t increases.

We can also make closely related predictions about the marginal impact of information used to form beliefs, i.e. dwelling characteristics and recent bills, on the probability of receiving an audit or making an installation as t grows. The marginal influence of this information should decrease the more observations a household has about their mean energy use. In our model, each of these factors increases the household's prior estimate of mean use m_t, which in turn affects whether the prior is above or below the ex ante usage threshold for an audit $\widetilde{m}$ or the ex post usage threshold for an installation $\bar{m}$. Priors with larger mean energy use make it more likely that the household will audit and install. However, the effect of θ and e_t on the prior at time t dampens as t grows. As the precision of the prior increases over time, θ and individual observations of e_t have smaller marginal effects on the household's estimate of m_t. In other words, the information contained in θ, and a particular e_t, has the biggest effect on expectations when the household is least certain about its average bills.

This is stated below as Lemma (2) (a proof is in appendix A.3):

Lemma 2 *Priors of mean energy use are increasing in home characteristics θ (e.g., age) and recent energy bills e_t, but at a decreasing rate across time:*

- $\frac{\partial m_t(\theta)}{\partial \theta} = m'_{t,\theta} > 0$,

- $\frac{\partial m_t(\theta)}{\partial e_t} = m'_{t,e_t} > 0,$

- $m'_{t,\theta}$ *and* m'_{t,e_t} *are declining in* t.

The longer a household delays an audit, the more precise their priors become and the less important observable characteristics of the home (θ) and any individual bill become in the updating process.

We can derive precise relationships on how the value of waiting changes by differentiating (8) with respect to θ or e_t.[15] We will focus our attention on the derivative with respect to θ because e_t operates on m_t in an analogous way, as shown in Lemma (2). In appendix A.4, we show that the partial derivative of (8) with respect to θ is given by

$$\frac{\partial E[V_t(audit)] - E[V_t(delay)]}{\partial \theta} = \left[\frac{p}{1-\delta}G(\bar{\alpha})(1-\widehat{\alpha})\Big(1-\delta+\delta H(\widetilde{m})\Big)+\delta h(\widetilde{m})(\tau+tr)A\right]m'_{t,\theta}$$
$$+\frac{\delta p}{1-\delta}\left[h(\widetilde{m})(\tau+tr)G(\bar{\alpha})\Big((1-\bar{\alpha})m_t-(1-\widehat{\alpha})\widetilde{m}\Big)+g(\bar{\alpha})\frac{\partial \bar{\alpha}}{\partial m_t}(1-\bar{\alpha})(1-H(\widetilde{m}))(m_t-\widehat{\mu}_H)\right]m'_{t,\theta}. \tag{9}$$

Equation (9) highlights that there are two competing effects dictating how the value of waiting changes based upon characteristics of a home. The first term is the direct effect of an increase in the prior on the likelihood of making an install if the household receives an audit. The first term has a positive effect on auditing in the current period rather than delaying. The second term is the indirect effect of an increase in the prior on the benefit of delaying in order to gain more information about the decision through updating. The second effect is negative. The negative effect is only large if the prior is far below the mean use at which making an install would be worth it. More precisely, if m_t is far below $\widetilde{m}$ or $\widehat{\mu}_H$ then the second term is large in magnitude.

An example is useful to gain intuition. Assume a household is unlikely to make an install because their prior mean is too low to justify it. In that case, a marginal increase in the prior mean makes the household more likely to delay in order to receive more

[15] In order to show rigorously how a longer history of bills impacts the audit versus delay decision, we must account for the effect of a change in beliefs about mean energy bills (m_t) on two parameters. The first is the installation savings threshold required to make an install ($\bar{\alpha}$). The second is the conditional mean of energy bills if beliefs are sufficiently high to justify an install.

information rather than incur an irreversible fixed audit or install cost. If the household's prior is already in the neighborhood of the level that would induce an audit, however, a marginal increase in the prior mean is likely to induce an audit in the current period rather than a delay. In other words, increases in the prior decrease the likelihood of delaying for those who are marginal to the decision, and increase the likelihood of delaying for those who were already predisposed to delay.

These two competing effects give rise to testable empirical predictions. The marginal effect of θ or e_t on the likelihood of auditing in the current period declines the longer the household has lived in the home.[16] Intuitively, the more time the household has lived in the home, the more precise their estimate of mean use. The characteristics that inform their original priors, as well as individual bill shocks, become less decision-relevant the less uncertainty the household has over its true mean. We state these results as a Remark:

Remark 3 *An increase in the prior estimate of mean use due to bill shocks e_t or home energy use characteristics θ*

- *increases the likelihood of auditing in the current period for households that are marginal to the decision;*
- *decreases the likelihood of auditing in the current period for households who already had a low audit likelihood;*
- *has a declining impact on audit likelihood as t grows.*

The model also makes empirical predictions about energy efficiency install behavior, conditional on an audit. Information from home characteristics θ and recent bill shocks e_t is also more relevant to the install decision the more uncertain the household is about its mean usage. Information has a larger marginal impact on the formation of the posterior if there are fewer historical observations from which to infer the mean, leading to the following Proposition:

Proposition 4 *Conditional on having received an audit, an increase in the prior estimate of mean use due to bill shocks e_t or home energy use characteristics θ*

[16] As shown in Lemma 2, both terms in equation (9) are multiplied by $m'_{t,\theta}$, which decline over time.

1. *increases the likelihood of making an installation following the audit;*

2. *has a declining impact on installation likelihood as t grows.*

Proposition 4 follows immediately from Lemma 2 (a proof is in appendix A.5). The testable implications are, first, homes with observable characteristics implying high energy use will make installs after an audit with higher probability than other homes, conditional on time in the home. Second, this difference in install rates will decrease over time in home.

In sum, there are several testable implications from the theoretical model that we can empirically investigate. The first set of predictions has to do with audit behavior and the second set of predictions analogously deals with install decisions. Each set of predictions relates to how incentives to audit and install change as households have different information about their home or a longer time series of information about their home.

3 Data

To explore whether the predictions stemming from our theoretical model are borne out in practice, we examine the energy efficiency investment decisions of households residing in a medium-sized MSA. During our study window (2011-2013), the local electric and gas utility provided residential customers with the opportunity to participate in a subsidized energy efficiency program. Specifically, households could pay a subsidized rate of \$50, rather than a retail rate of \$150, and receive an energy audit. Households that scheduled an audit had an expert with specialized training and equipment visit their home and received advice on which, if any, investments could be made to meaningfully improve the energy efficiency of their homes. Households choosing to make energy efficiency upgrades based upon that advice received up to \$500 from local power providers in rebates. To receive the installation rebate, a household must have first had an audit.

To test the predictions from the theoretical model, we examine how the likelihood of participating in the audit program and making subsequent energy efficiency upgrades varies with the characteristics of a home and the time spent living in the house. To do so, we combine three unique datasets. The first dataset we use records the addresses of every household that scheduled an audit from 2011-2013. The dataset also includes all install decisions. For both audits and installs, the dataset records the date of the audit and the installation.[17] The second dataset is county assessor data at the address level. This data includes characteristics of each home such as square footage, year built, type of heating, number of stories, etc. We leverage this data to determine how audit probability varies with home characteristics. For premises that were sold at any point during our sample period, we also observe the date of sale.[18] The third and final dataset is address-level, monthly electricity and natural gas billing data spanning 2011-2013. The data includes aggregate household electricity and gas consumption as well as the billing period start and end dates.[19]

In total, there are 150,658 unique premises in our billing data sample. Ideally, we would be able to focus exclusively on the owner-occupied premises in our sample. Unfortunately, that information is not available. In order to exclude premises that are highly likely to be rental units, we drop each premise that has "Unit" or "Apt" in the address.[20] Ultimately, this leaves us with a final sample of 88,791 unique premises. Consistent with other audit studies, we observe low take-up rate of audits: 2,573 audits over three years.

Table 1 summarizes the premises in our sample. The homes are divided into those that do not receive an audit during 2011-2013 and those that do receive an audit. The summary statistics highlight that the audited homes are, on average, older and larger. The mover indicator is equal to one if a premise is ever sold to a new owner during the 2011-2013 period. Importantly, Table 1 highlights that there is a positive correlation between a premise being purchased and audited. In particular, during the period spanning 2011-2013, 8% of the premises were sold. Among these premises, 4% received an audit during our sample period – compared to only 2% among the homes that were not sold during our

[17] The data was shared under a privacy agreement directly by the auditing agency.

[18] This data was publicly available from the county assessor in our study footprint.

[19] The frequency of household billing data is remarkably stable with a billing period every 30 or 31 days for all households. The exceptions are often very short bills followed by a gap then another bill over a short period. In conversations with the utility, these are almost always billing interruptions due to the sale of homes or changes in the renter of rental units. Exploratory analysis comparing the bill dates to the sale dates in county assessor data confirms this. This data was also shared under a privacy agreement directly by the utility.

[20] Combining the program participation data with the assessor data requires matching the premises based on their addresses. In some cases, the form in which the addresses enter differ across the assessors data and the utility data. In cases where an address match does not exist, we use a text matching algorithm to match premises across the two datasets. Ultimately, we error on the conservative side and drop premises which do not have a clear address match across the two samples.

sample period.[21] In the following section, we explore this pattern in much greater depth. In particular, we first explore what affects if, and when, households elect to receive an audit. Next, we explore what impacts households' subsequent decisions to make energy efficiency upgrades to their homes.

Recall from the theoretical model, our key assumption is that households do not know how much energy they will consume in a home prior to living in it for an extended period of time. While we do not have anyway to observe or measure peoples' beliefs, we do present evidence in Appendix B highlighting that energy use varies dramatically across premises, even after conditioning on observable characteristics. Again, this opens the door to the possibility that some homeowners could move into homes that appear to be "energy hogs" based upon observables and make an install based upon expected savings only to discover the home was actually energy efficient and the install was a mistake – a possibility that we will explore in greater depth following the empirical tests of the theoretical model.

4 Testing Theoretical Predictions

4.1 Are Audits More Likely Immediately After Sales?

Recall, our model predicts that a household will choose a costly audit if they think their mean energy usage (μ) is large enough, and the share of energy consumption that would remain after investing in energy efficiency upgrades (α) is small enough, to make it likely that upgrades would ultimately be beneficial. However, both μ and α are unknown parameters. When a household moves into a home, their expectation of their mean energy usage is a function solely of observable characteristics (e.g., the home's age, condition). This prior is updated as the household spends time in their new home (e.g., receiving monthly energy bills and experiencing the comfort levels during the winter and summer).

From this simple framework, there are several predictions surrounding if and when households elect to receive energy audits. First, immediately after moving into a home, we would expect to see a mass of households select to receive an IHEA. This mass would be comprised of the households that had initial priors for μ that were sufficiently high to

[21] Data restrictions prevent us from observing whether homes had received audits prior to our sample period.

justify an immediate audit. Therefore, empirically, the first question is, do we see this mass of "immediate auditors"?

Focusing on the premises that were audited and sold at some point during our 3-year sample period, we calculate the number of days between the recorded sales date and the date the IHEA was scheduled. Figure 2 displays the histogram of the number of months between the audit date and the sale date. Consistent with our theoretical model, there is a clear spike up in audits occurring during the first month following the home sale.[22]

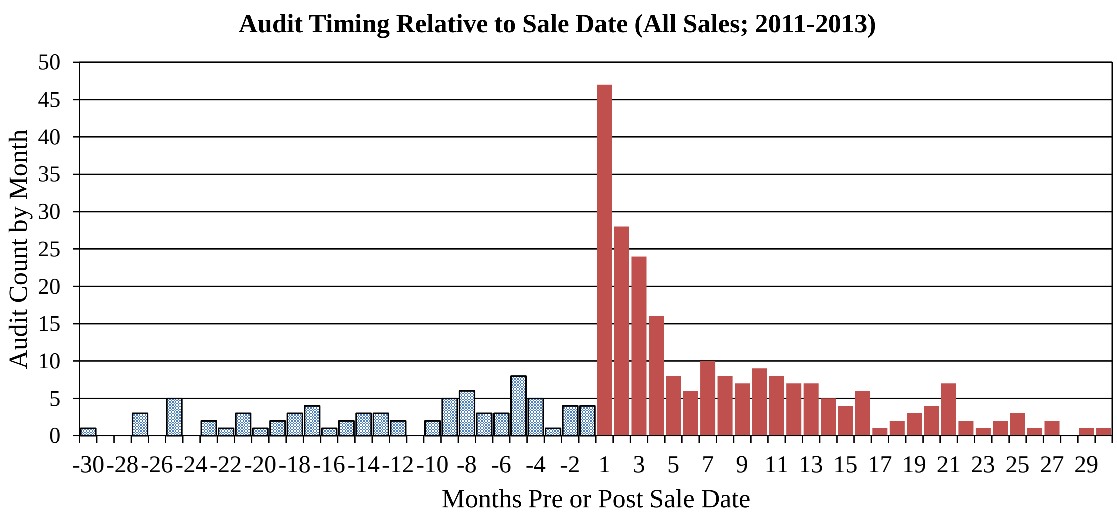

Figure 2: Histogram of number of Audits relative to move date for homes where an audit occurs and the home is sold between 2011 and 2013.

Our model also suggests that the premises audited immediately will have different household observables (e.g., home age) relative non-auditing households. Households opting to receive audits immediately after moving should have high initial priors for their premises' mean energy usage (μ) relative non-auditing movers. Given that the new homeowners have no experience in their homes, their initial expectations of μ are based solely on observed characteristics. One potentially relevant characteristic – that we as researchers

[22]To ensure that we observe a full year pre and post-sale date during our sample period, Figure A3 recreates the same histogram focusing on the audited premises that were sold during the middle year of our sample (2012). Again, there is a large spike in audits immediately following the sale. Similarly, A4 displays the probability (frequency) a home is audited during a given month after moving in, conditional on not being audited in the prior months. Again, the figure displays the probability of being audited falls with time.

can also readily observe – is the year the home was constructed. In general, older homes are expected to be less energy efficient. In part, homes depreciate with age (e.g., insulation becomes less effective over time; as homes settle, cracks and air leaks develop). Moreover, as building codes become stricter and technology improves, energy efficiency typically increases across vintages.

To explore whether homes audited immediately after being sold have different observable characteristics than homes audited later, we compare the observables for homes audited within 30 days of being purchased to the homes that are purchased and audited more than 30 days after being audited. Table 2 displays the summary statistics for these two subsets of sold and audited permises. Focusing on the year of construction, we see that among the 48 premises that were audited during the first month post-sale, the average (median) year of construction was 1964 (1963). In contrast, among the 183 premises that were audited more than 30 days after the sale date, the average (median) year of construction was 1968 (1967). Given the fairly limited sample size, these differences in means are not statistically significant. However, the pattern is consistent with the prediction that the premises audited immediately are older, and would therefore plausibly come with higher initial expectations for μ.

4.2 What Explains the Delayed Audits?

It is important to note that there are potential alternative explanations that could result in a similar mass of immediate auditors. One notable possibility is that there may be convenience costs incurred by having an audit, and subsequent energy efficiency upgrades, performed. It is certainly possible that these convenience costs could discontinuously increase after the first several weeks following a home sale. For example, onces new homeowners move their belongings into their new home, it may become more challenging to have and inspection and upgrades performed on the home.[23]

However, a simple discontinuous increase in the private costs incurred by having an audit or upgrades performed would not be able to explain the slow decay in the frequency of audits several months after a premise is sold (i.e. the fat right tail of Figure 2). Instead, the

[23]This assumption could be included in our model framework simply as a discontinuous increase in the fixed audit and upgrade costs, A and F, following the first period.

steady decline in audit frequency post-sale is potentially consistent with another prediction stemming from our theoretical model. In particular, new homeowners that do not elect to receive an IHEA immediately after purchasing their home begin to receive information that results in updates to their expectation of μ. If this new information results in them updating their expectation of μ upwards over time, then we would expect to see additional households select to receive audits. Importantly, our model predicts that, over time the households' expectation of μ become more precise. As a result, the initial information (e.g., the first few energy bills) will be the most influential in terms of moving the households' priors surrounding the mean energy usage. Therefore, consistent with Figure 2's steadily declining right tail, we would expect to see fewer and fewer homes electing to receive audits as the months post-sale increase and meaningful movements in $E[\mu]$ become rare.

To explore whether there is evidence that households are updating their expectations of μ over time, and subsequently deciding to receive audits, we explore whether the timing of the audits that occur beyond one month after the sale date coincide with periods of high energy consumption. To do so, we focus on the set of households that receive an audit at least one month after purchasing their home. Using the recorded monthly electricity and natural gas consumption from each individual household, we are able to explore how the likelihood of receiving an IHEA during a given month responds to the contemporaneous and lagged energy consumption using the following linear probability model:[24]

$$\text{Audit}_{i,t} = \alpha_i + \gamma_m + \beta_1 \cdot \text{Elec}_{i,t} + \beta_2 \cdot \text{Gas}_{i,t} + \theta_1 \cdot \text{Elec}_{i,t-1} + \theta_2 \cdot \text{Gas}_{i,t-1} + \varepsilon_{i,t}. \quad (10)$$

In the model specified by Eq. 10 above, $\text{Audit}_{i,t}$ is an indicator variable which equals one during the month t in which household i receives an energy audit.[25] $\text{Elec}_{i,t}$ and $\text{Gas}_{i,t}$ represent the average daily electricity (kWh) and natural gas (therms) consumed by household i during month t's billing cycle.[26] $\text{Elec}_{i,t-1}$ and $\text{Gas}_{i,t-1}$ are the average daily

[24]For this exercise, it is important to include the households' monthly electricity and natural gas consumption. While electricity use typically peaks in the summer in the study region, natural gas consumption tends to peak in the winter. Suppose, for example, households were driven to receive audits following high gas bills. If we focused exclusively on how the likelihood of an audit responds to electricity use, then we would find audits are more likely to occur during low electricity consumption months. By restricting our sample to households that have observed electricity and gas consumption and receive an audit beyond one month after the sale date, we are left with 86 households.

[25]Once a household is audited, we drop the remaining monthly observations from the household.

[26]The billing cycle start and end dates vary across households, and therefore, do not necessarily coincide

consumption levels during the preceding bill cycle. Household fixed effects are included to control for the fact that there is substantial heterogeneity in energy use across households. In addition, we estimate the model with and without monthly fixed effects.[27]

The key coefficients of interest from Eq. 10 are $\{\beta_1, \beta_2\}$ and $\{\theta_1, \theta_2\}$. If households update their expectations of μ as they receive energy bills, then we would expect to see an increase (decrease) in the likelihood of an audit occurring during a month preceded by high (low) levels of electricity or natural gas consumption. This would imply that θ_1 and θ_2 are positive. An alternative possibility is that households elect to receive an audit during a month in which their energy consumption is abnormally high – perhaps due to a real-time awareness of their energy usage or perhaps due to uncomfortable living conditions. In this case, β_1 and β_2 would be positive.

Table 3 presents estimates from the general model specified by Eq. 10. The first two columns present estimates in which we restrict θ_1 and θ_2 – the coefficients on the lagged energy consumption – to be zero. Column one and two report the estimates of β_1 and β_2 without and with monthly fixed effects, respectively. Across both specifications, we see there are positive point estimates for both coefficients – suggesting that audits are more likely during months with high energy consumption. In columns three and four, the estimates of Eq. 10 are presented with the lagged energy consumption included in the model. Across both specifications (i.e. with and without monthly fixed effects), there is clear evidence that the likelihood of a household choosing to receive an audit increases during a month that follows a period of high electricity or natural gas consumption. This suggests that households indeed display evidence of updating behavior consistent with our theoretical model. Specifically, among the households that do not immediately receive audits after moving into their new premises, high realizations of energy consumption can subsequently nudge them towards receiving an audit.

Again, our theoretical model predicts that, over time, as households observe a longer series of bills, their expectations of their mean energy usage μ becomes more precise.

with the calendar months. To create the monthly fixed effects, we assign each billing cycle to the calendar month in which the majority of the cycle's days occurred.

[27]While including the monthly FE will control for seasonal patterns in audit likelihood that are not driven by energy consumption, they may simply end up sweeping away the seasonal variation in energy consumption that may explain the audit take-up.

While observations of high or low energy bills in the months after a household moves into a new home may meaningfully move their prior surrounding the mean energy consumption, observed bills would have little impact on the expectation of μ once the household has resided in the home for a longer period of time. Therefore, our theoretical model predicts that, among the households that have not recently moved into their homes, contemporaneous or lagged energy bills will not meaningfully predict whether the household elects to receive an audit. To test whether this is observed in our setting, we reestimate the model specified by Eq. 10, this time focusing exclusively on the households that are audited but had moved into their homes prior to 2011, before our sample period begins. Columns five and six of Table 3 present the estimates of Eq. 10 without the lagged electricity and gas consumption – once without monthly fixed effects and then again with. Columns seven and eight present the estimates with the lagged energy usage included. In contrast to the estimates from Columns one through four, there is no consistent evidence that past energy consumption meaningfully impacts the likelihood of receiving an audit among the non-movers.

4.3 Differences Between Audited and Non-Audited Premises

There are also testable implications from the theoretical model surrounding not just when, but if a given household will elect to receive an audit. Recall, our model predicts that an audit occurs if a household's expectation of the mean energy use, μ, is sufficiently high. As we noted earlier, the expectation of μ will be, in part, a function of the premise's observable characteristics – such as the year the home was constructed. Consistent with the theory that observables can affect the likelihood of an audit, Table 1 highlights that, across the full set of premises in our sample, homes that were audited during our three year sample are indeed significantly older.

Importantly, the theoretical model predicts a more nuanced relationship between a home's age and the likelihood of it being audited. While an observable characteristic like the year of construction can play an important role in determining the $E[\mu]$ among households that recently moved into a premise, households that have resided in their homes for many years will have formed fairly precise expectations of μ based on past consumption as opposed to simple observables such as age. This suggests that the relationship between the likelihood of an audit occurring during our sample and the age of a home should be much stronger among households that have recently moved as opposed to those that have been in their homes for multiple years.

To explore whether the relationship between audit likelihood and year of construction vary across movers and non-movers, we focus on all of the premises in our sample constructed between 1950 and 2000.[28] We classify a household as a mover if an audit was received anytime after the move-in date during our three year sample. Figure 3 plots locally smoothed polynomials displaying the relationship between the likelihood of an audit occurring and the year of construction for movers and non-movers. Consistent with the theoretical model's prediction, there is a clear negative relationship between year of construction and audit likelihood among movers. This relationship effectively disappears among the non-movers.

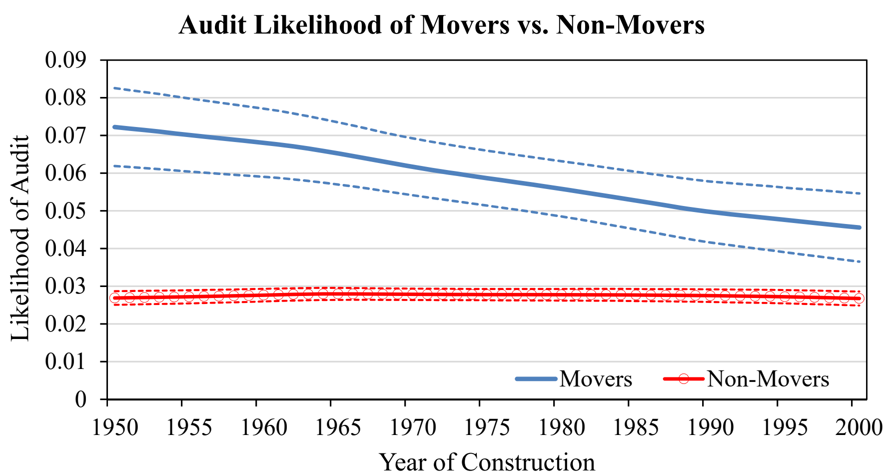

Figure 3: Likelihood of audit by home construction year conditional on a home being sold (Movers) versus not sold (Non-Movers) in our sample.

Of course, a number of factors may be correlated with home age. Therefore, the negative correlation between year built and audit likelihood among recent movers could

[28]Prior to 1950, the year of construction variable is coarsely aggregated by decade – without any information on whether the year of construction is rounded up or down. Therefore, we do not include premises with year of construction reported prior to 1950. In addition, audits, and subsequently upgrades, are quite rare among the premises that are less than 10 years old. Therefore, we do not focus on newly constructed homes.

instead be caused by a link between audit likelihood and, say, the income levels of the neighborhoods with homes of different vintages. To further test whether the year of construction indeed affects the probability of a household electing to receive an audit – and whether this relationship differs across movers and non-movers – we examine how the probability of an audit varies with the year of construction while controlling for a wide range of other observable premise characteristics. To do so, we estimate the following linear probability model:

$$\text{Audit}_i = \alpha + \beta \cdot \text{Year}_i + \theta \cdot \text{Year}_i \cdot \text{Mover}_i + \boldsymbol{\phi} \cdot \boldsymbol{X}_i + \varepsilon_i. \tag{11}$$

In the model specified by Eq. 11, Audit_i is an indicator variable which equals one for each premise i that is audited at any point during our three year sample.[29] Similarly, Mover_i is an indicator variable which equals one for each premise that is sold during our three year sample. Year_i is equal to the year the home was constructed. Finally, to control for other characteristics that may be correlated with the year of construction, the vector $\boldsymbol{X}_i$ includes a fully saturated set of 260 indicator variables separating houses into groups based on the Mover_i indicator, the number of bedrooms $(1, 2, ..., 5+)$, whether the house is single versus multi-story, and the square footage (thirteen bins ranging from < 800 square feet to $> 3,200$ square feet). In addition, to control for unobserved differences across neighborhoods, we include a set of 35 postal code fixed effects as well as the interaction between the postal code effects and the mover indicator.

In Eq. 11, β captures how the likelihood of an audit occurring at non-mover premises changes, on average, if the home were one year newer (i.e. the year of construction increases by one year). Similarly, $\beta + \theta$ represents how the audit likelihood changes, strictly among the homes that are sold, as the premise age falls. Table 4 presents the point estimates of β and θ. The first column does not include the observable controls ($\boldsymbol{X}_i$). The second column includes only the premise characteristics. Finally, column three includes controls for the premise characteristics as well as the zip code fixed effects. Across each specification, a clear pattern emerges. As our theory model predicts, among the homes that are sold during

[29]Note, in contrast to the model specified by Eq. 10, which explored the timing of audits, the audit indicator now only varies across premises, not across time.

our sample period, the likelihood of an audit occurring falls as the year of construction increases (i.e. as the age of the home falls). The point estimates suggest that the likelihood of observing a premise sold and subsequently audited during our sample period falls by approximately 1 percentage point if the age were to fall by 10 years, all else equal. In contrast, among households that resided in the same premise through the full three-year sample period, the likelihood of observing an audit only falls by 0.1 percentage points if the age of a home fell by 10 years.

Consistent with our theoretical model's prediction, the above results provide suggestive evidence that new homeowners use observable home characteristics (e.g., home age) to inform their decision on whether to receive an audit or not. Our theory model also predicts that this relationship should weaken over time. That is, as homeowners spend more time in their new homes and gain more information surrounding their home's true energy requirements, the importance of observables (e.g., home age) in determining whether to receive an audit or not should decline. To explore whether this prediction holds true in the data, we explore how the likelihood of an audit varies with a home's age among different subsets of movers. Figure **??** again plots a locally smoothed polynomial (the upper solid line) displaying the relationship between the likelihood of an audit occurring and the year of construction for all movers. In addition, the figure also displays the same relationship among the movers that do not receive an audit within the first two months of residing in the home and among those that are not audited with in the first four months of residing in their new home. Consistent with the theory model, the slope declines with time since move-in.

To test whether the relationship between home age and audit-likelihood indeed decays over time, we re-estimate the model specified by Eq. 11 excluding the homes that are audited within the first two months of moving (results reported in column four of Table 4). Consistent with Figure **??**, the impact of the year of construction on the likelihood of audits among the movers that do not audit immediately declines (the coefficient on Year Built$\times$Mover falls to -0.004).

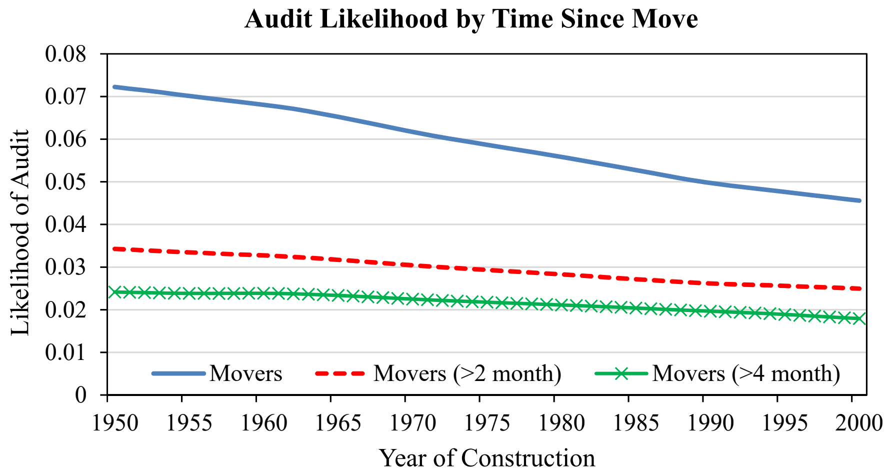

Figure 4: Notes

4.4 Installation Decision

The final set of testable predictions from the theoretical model focus on households' decisions to make energy efficiency upgrades following the receipt of an IHEA. Recall from the model, after a household opts to receive an audit, they will be fully informed about α – the share of energy consumption that would remain if they were to make the recommended upgrades. Armed with the knowledge of α, a household must then decide whether to perform the energy efficiency upgrades. If the household's expectation of μ – their mean energy usage absent performing any upgrades – is sufficiently high, then they would opt to perform the upgrades. However, if $E[\mu]$ is not large enough given the realization of α, then no upgrades would be performed.

Therefore, the first installation-related question is the following – do we observe households electing not to make energy efficiency upgrades after receiving an audit? Table A1 summarizes the frequency with which different types of energy efficiency upgrades were performed following the audit. The frequencies are reported separately for audits that occurred among premises sold during our sample period (i.e. movers) as well as those that were not sold during out sample. The table highlights that energy efficiency upgrades were ultimately performed in 63% of the 308 homes that were sold and audited. Similarly, upgrades were performed in 62% of the audited homes that were not sold during our three year sample. Consistent with the theoretical model, just over one third of the audited homes elect not to perform any upgrades. Interestingly, this share does differ significantly across movers and non-movers in our sample. This suggests that individuals' beliefs over their mean levels of energy use do not systematically differ across movers and non-movers. For this to be the case, households that recently move, and therefore base their expected energy use heavily off the observed characteristics of the home, must be forming unbiased priors of their mean energy usage.

Beyond simply suggesting that some audited households would not perform energy efficiency upgrades, the theoretical model also predicts that the likelihood of post-audit upgrades being performed would be a function of observable premise characteristics. In particular, premise characteristics that imply a higher level of expected energy use (e.g., an older home) can affect the install decision for households who have recently moved into a home and are still basing their beliefs about μ on these observables as opposed to a long time series of information gained through energy bills. In contrast, among households that have resided in their homes for longer periods of time, observable characteristics may have little impact on the likelihood of energy efficiency upgrades being performed post-audit.

To explore whether the likelihood of energy efficiency upgrades occurring varies with the age of a home, we explore how the installation frequency varies among audited homes constructed between 1950 and 2000. Figure 5 plots locally smoothed polynomials displaying the relationship between the frequency of any upgrade occurring and the year of construction for movers and non-movers. Consistent with the theoretical model's prediction, there is a negative relationship between the likelihood of an upgrade being performed and the age of the premise among homes sold during our sample period. This negative relationship, while still visible, is weaker among the audited homes that are not sold during our sample period. Figure A5 in the appendix displays similar locally smoothed polynomials comparing the likelihood of different types of energy efficiency upgrades being performed and home vintage. The same patterns emerge – premise age appears to have a stronger positive impact on the likelihood of upgrades being performed among the premises that are recently sold.

Again, the year of construction can certainly be correlated with other characteristics that may affect the likelihood of upgrades being performed. To further test how premise

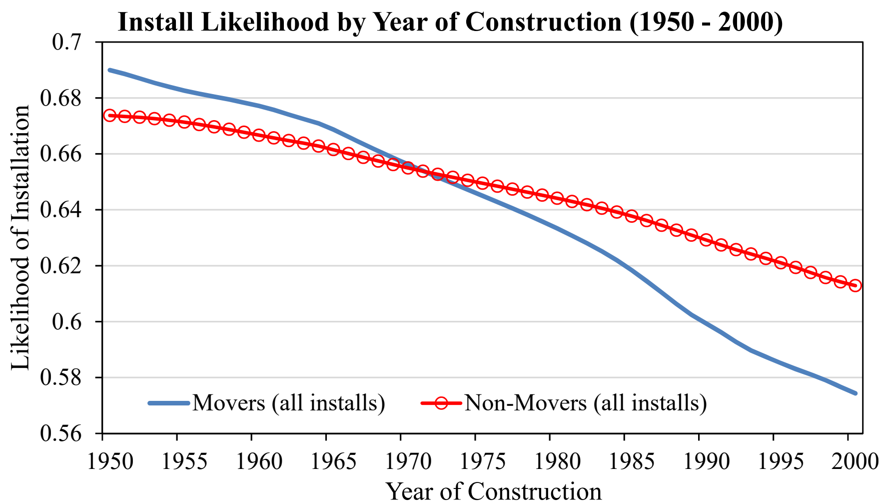

Figure 5: Install likelihood conditional on having an audit for home that are sold (Movers) versus not sold (Non-Movers) by year of construction. The Figure shows movers are much less likely to make installations in newer homes conditional on having an audit and more likely to make installs in older homes.

age and upgrade frequencies are related, we estimate the following linear probability model focusing exclusively on the premises that are audited during our sample period:

$$\text{Install}_i = \alpha + \beta \cdot \text{Year}_i + \boldsymbol{\phi} \cdot \boldsymbol{X}_i + \varepsilon_i. \tag{12}$$

In the model specified by Eq. 12, Install_i is an indicator variable which equals one for each premise i that performs any energy efficiency upgrade following their IHEA. Year_i is again equal to the year the home was constructed. Finally, to control for other characteristics that may be correlated with the year of construction, the vector $\boldsymbol{X}_i$ includes indicator variables reflecting the number of bedrooms $(1, 2, ..., 5+)$, whether the house is single versus multi-story, and the square footage (thirteen bins ranging from $<$ 800 square feet to $>$ 3,200 square feet). In addition, to control for unobserved differences across neighborhoods, we include a set of 35 postal code fixed effects.

The model specified by Eq. 12 is estimated separately for the 243 audited homes that are sold during our sample period (i.e. Movers) and for the 1,387 audited premises that are not sold during our sample (i.e. Non-Movers). Columns one through four of Table 12 present the point estimates of β with and without the observable controls. Consistent with the pattern displayed in Figure 5, the likelihood of upgrades occurring falls more rapidly with the year of construction among the Movers compared to the Non-Movers. Indeed, with the full set of premise characteristic and zip code controls included, the relationship between year of construction and upgrade likelihood is insignificant among the Non-Mover and negative and statistically significant among the Movers. Columns five and six of Table 5 include point estimates of the difference between β from Eq. 12 for Movers and Non-Movers. While the difference between the Movers' and Non-Movers' β's are not found to be statistically different, the pattern is consistent with the theoretical prediction. Specifically, recent movers' decisions to perform upgrades appears to be more heavily influenced by observables such as home age.

5 Investment Mistakes and Policy Implications

The preceding analysis reveals that the predictions stemming from the theoretical model are largely borne out in the data. In this section, we return to the theoretical model to garner insights surrounding the efficacy of programs subsidizing energy efficiency audits and upgrades. We first highlight that, given the uncertainty surrounding the true payoffs from investing in energy efficiency, households may make an install even though they should not. We then demonstrate that these investment "mistakes" can be exacerbated by typical energy efficiency audit and install subsidy programs. While it is well known that in this type of setting, a large share of participating households may be inframarginal or "non-additional" to the program, our analysis shows that even the households that are technically "additional" to the program may not be economically efficient participants.

From the theoretical model, we can begin by dividing households into four categories: (1) households that correctly make energy efficiency upgrades, (2) households that correctly choose not to make upgrades, (3) households that delay making upgrades that would be economically efficient, and (4) households that make upgrades which are inefficient and don't pass a full information benefit-cost test. Which group a given household falls into depends crucially on their prior beliefs about their energy usage and the true mean usage. Figure 6 plots an example compound distribution of prior beliefs and true mean usage. The solid black lines represent a given $\bar{m}$, the *ex post usage threshold*, which is the cutoff mean usage above which making an install is efficient if the household knows the true savings α and their true mean μ. If the true mean is above the threshold, the household should make an install. However, because of uncertainty in the true mean, households will only make an install if their prior belief (m_t) about their mean is above this threshold.

From Figure 6, we can characterize four types of install decisions at any given time, and discuss how an install subsidy or an audit program alters the probability of falling into each category:

1. $\mu > \bar{m},\ m_t > \bar{m}$, optimal to install: the household invests in energy efficiency and it is efficient to do so. This is represented in Figure 6 by the area under the compound distribution in the dark grey box in the upper right corner labeled "Correctly Install".

2. $\mu < \bar{m},\ m_t < \bar{m}$, optimal not to install: the household does not invest in energy efficiency, which is economically efficient. This is represented in Figure 6 by the unshaded region in the lower left.

3. $\mu > \bar{m},\ m_t < \bar{m}$, delay mistake: the household does not yet invest in energy efficiency when it should. This is represented in Figure 6 by the medium gray rectangle in the upper left labeled "Don't Install, But Should". In our model, learning over time causes the prior to converge to the true mean so that as these households update their priors they will eventually move in to the dark gray "Correctly Install" region. In Figure 6 this means that eventually the mass of the compound distribution is concentrated on the 45 degree line as uncertainty about the true mean usage declines.

4. $\mu < \bar{m},\ m_t > \bar{m}$, investment mistake: the household invests in energy efficiency when it should not. The households that fall into the light gray box in the lower right of Figure 6 make this irreversible investment mistake. They make an install right away because their prior belief is above the threshold. Yet as their beliefs are updated over time and converge to their true mean usage, these households eventually learn that their investment was inefficient for them.[30]

[30] One could more specifically quantify how the share of households in each category responds to policy by parameterizing a particular functional form for the compound distribution. The probability that a household makes a mistake (household type 4) is given by the joint probability

$$Pr(\mu < \bar{m},\ m_t > \bar{m}) = \int_{-\infty}^{\bar{m}} h(\mu|m)d\mu \int_{\bar{m}}^{\infty} dM$$

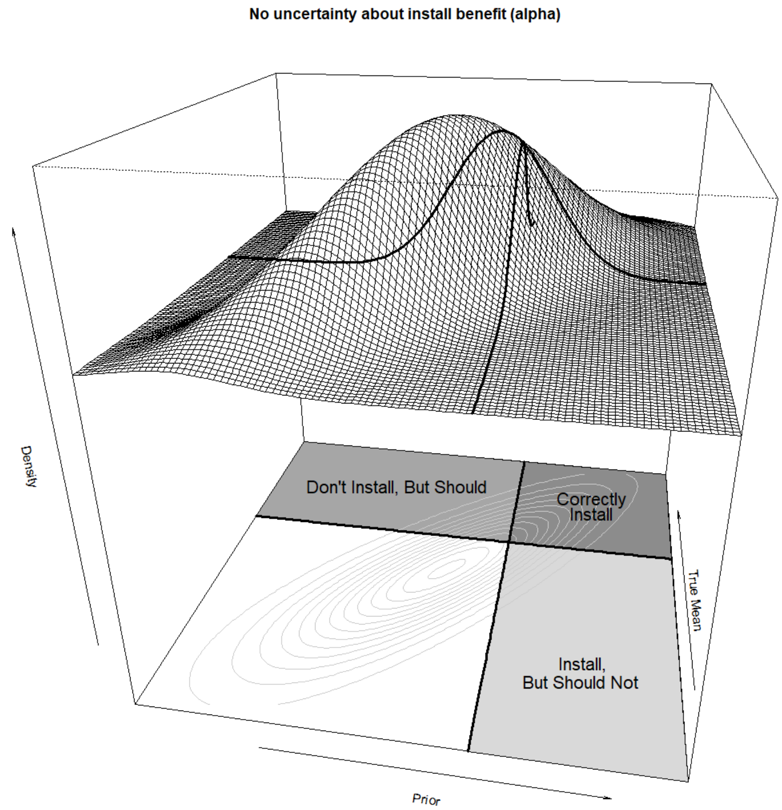

Figure 6: Notes: This figure delineates four categories of households. The share of households that install are those with a prior above the cutoff, given by the area under the compound distribution to the right of the vertical black line ("Install, But Should Not" and "Correctly Install"). The share of households that should install are those with a true mean usage above the cutoff, given by the area under the compound distribution above the horizontal solid black line ("Don't Install, But Should" and "Correctly Install"). The unshaded area in the bottom left are those who make an optimal decision not to install.

To begin, we focus on how the decision to make an installation is affected by an upgrade subsidy. To do so, we focus on the case where the savings parameter (α) is known – i.e. we assume an audit has already occurred. By reducing the private cost of making an installation, the upgrade subsidy reduces the installation threshold $\bar{m}$, which we depict in Figure 7 as a reduction in the installation threshold from $\bar{m}_0$ to $\bar{m}_1$. Consequently, a subset

where the expression is derived using Bayes Rule. Comparative statics or numerical sensitivities using this expression can be calculated in a straightforward way in order to characterize how an audit mediates or moderates the effects of an install subsidy, or show that as uncertainty in μ declines the share of households making mistakes also declines.

of the type (2) households, for whom it was optimal not to install without the subsidy, now make an installation. Of these, only the households in the thickly cross-hatched area labeled “Target Additionality” become type (1) households for whom it is now optimal to install. These are the households the program intends to target. Models without uncertainty and learning typically have these households in mind as being marginal or additional to the program.

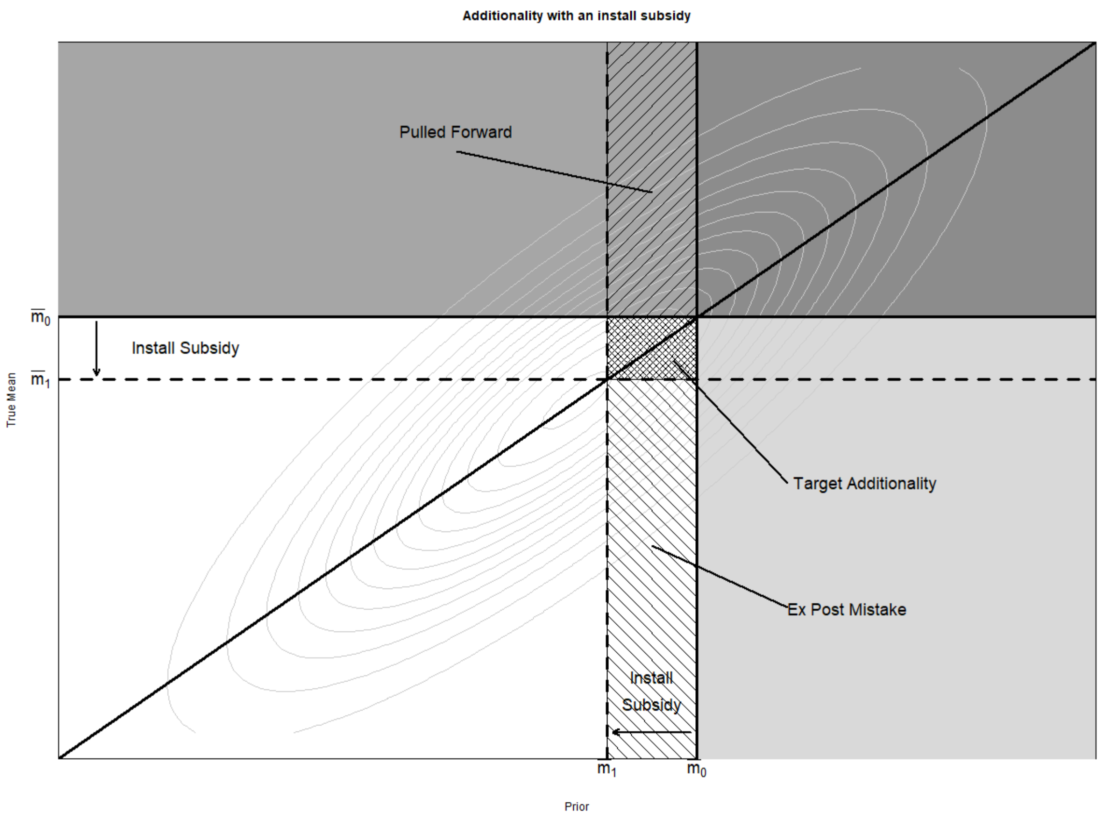

Figure 7: Notes: A subsidy reduces the critical value $\bar{m}$ for households to make an install. The share of households that are “additional” to a subsidy policy are the area under the compound distribution between the vertical dashed black line ($\bar{m}_1$) and the vertical solid black line ($\bar{m}_0$). However, only those in the cross-hatched area labeled “Target Additionality” are economically efficient participants. The rest have a true mean below the critical value (“Ex Post Mistake”) or would have eventually made an install without the subsidy (“Pulled Forward”). The figure depicts the case when α is known.

As the figure shows, however, a potentially significant share of the technically additional households may not be economically efficient. For example, the remaining households that did not install without a subsidy - those in the white, lightly hatched area below the "Target Additionality" group - are induced to make an irreversible investment mistake because of the subsidy, becoming type (4) households by the reduction in $\bar{m}$, given uncertainty over μ. In addition, some of the type (3) households are "pulled forward" into making an installation earlier than they otherwise would have. These are the households in the medium gray, lightly hatched area above the "Target Additionality" region. It is optimal for these households to make an installation, but they will eventually do so in the absence of a subsidy as their prior m_t converges to $\mu > \bar{m}$ over time. Ultimately, the economic efficiency gains stemming from this pull-forward effect depend on the discount rate. Although all three new adopters are technically "additional" to the policy in that the subsidy caused their participation, the "Target Additionality" group are the only unambiguously economically efficient adopters.

Ultimately, the "Target Additionality" group may be a large or small share of the households induced to install by the subsidy, depending on the shape of the compound distribution of prior beliefs and true means. Importantly, among households that have just moved into a home, we would expect the correlation between the prior beliefs and the true means to be the lowest. Consequently, among these recent movers, the share of participants in the "Target Additionality" region would be the smallest. As households update their priors, uncertainty about their true mean usage declines and the compound distribution converges to the 45 degree line. Once there is no remaining uncertainty, all of the households that are induced to install by the subsidy are part of the "Target Additionality" group. This suggests that likelihood of inducing investment mistakes (i.e. Type 4 households) is the greatest when households participate in energy efficiency subsidy programs immediately after moving into a new home.

Focusing next on the impact of subsidizing audits, we find a very similar pattern. An audit that reduces uncertainty about the savings parameter α without reducing uncertainty about true mean usage can exacerbate the occurrence of inefficient uptake by inducing more installs without fully informing the investment decision. In appendix A.6, we show that the *ex post usage threshold* $\bar{m}$ – the threshold belief when α is known – is below the threshold belief when α is uncertain (the *ex ante usage threshold* $\widehat{\bar{m}}$). This situation is depicted in Figure 8. In Figure 8, $\widehat{\bar{m}}$ is depicted with the dashed black line, and the full information $\bar{m}$ is depicted with the solid black lines. Audits induce more installations because they reduce the installation threshold through the resolution of uncertainty in α. Doing so may again "pull forward" a share of households that were delaying their installation when it was optimal (the shaded region between the solid and dashed lines). These households would have eventually become inframarginal to the policy through updating their beliefs over time, so the economic efficiency gains of this pull-forward effect may again be small. Because of uncertainty in true mean usage, however, the audit also increases the number of households who make an irreversible investment mistake which increases economic inefficiency.

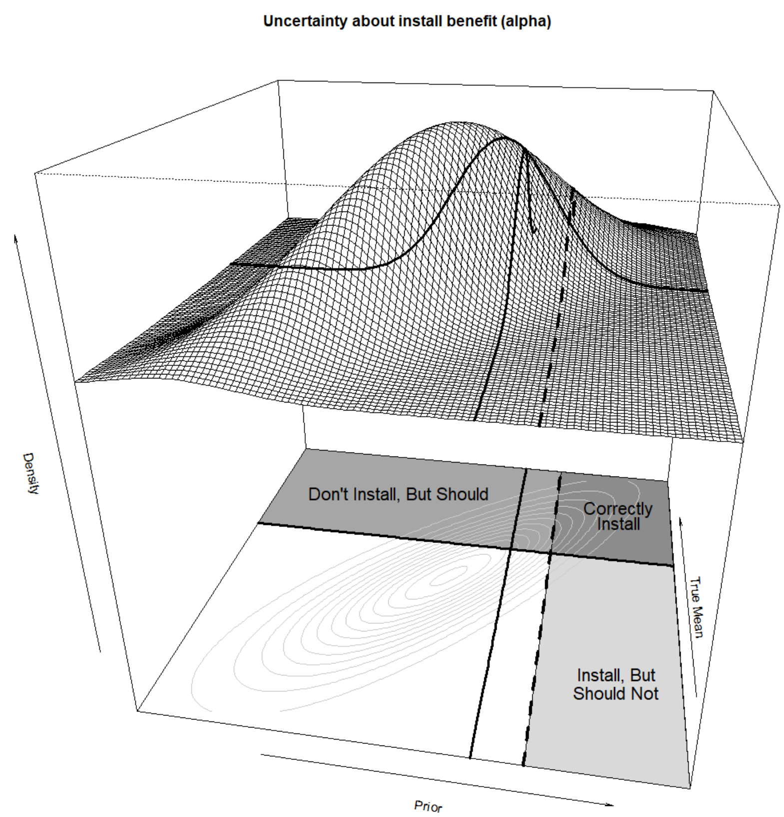

Figure 8: Notes: The share of households that install when α is uncertain is the area under the compound distribution to the right of the dashed black line ("Install, But Should Not" and "Correctly Install"). If uncertainty in α is removed, this cutoff moves to the left, to the solid black line. These additional participants are not necessarily economically efficient. The share of households that should install is the area under the compound distribution above the horizontal solid black line ("Don't Install, But Should" and "Correctly Install").

The insights provided by the theoretical model first suggest that the social benefits provided by subsidizing energy efficiency upgrades may be far smaller than previously thought. Recall, existing studies demonstrate that only a fraction of subsidized energy efficiency upgrades represent additional investments (e.g., less than 50% in the setting explored by Boomhower and Davis (2014)). Our analysis highlights that many of these additional investments are not truly additional, but rather simply pulled forward in time. Even more troubling, given the existence of uncertainty in the decision making process, we highlight that many of the additional participants may be making inefficient investments – i.e. upgrades with costs larger than the stream of social benefits they will provide. Quantifying the share of additional investments that are either pulled-forward or are truly mistakes would clearly be a challenging task. To do so, one would need information on the joint distribution between the actual returns to energy efficiency investments as well as the households' beliefs regarding the benefits – which would include not only the benefits resulting from bill savings, but also the comfort benefits which we have abstracted from in our simple model. While it is beyond the scope of the present study, our results suggest that it would be particularly useful for future work to explore the relationship between households' beliefs and the true returns to energy efficiency.

More generally, our analysis highlights that efficiency gains could be achieved by eliminating the uncertainty households face when deciding to make energy efficiency upgrades. Importantly, as Figure 7 highlights, it is not sufficient to simply remove the uncertainty surrounding the share of energy use that could be avoided (α). It is also important to ensure that households have a precise understanding of their baseline energy usage (μ). While households will ultimately update their priors and learn what their true μ is after living in their homes for a period of time, households that have just moved in will still be confronted with uncertainty in μ. Given our theoretical prediction – and empirical observation – that households are most likely to participate in energy efficiency programs at the time they move into a home, this suggest that a program targeting recent movers with information about both α and μ could substantially reduce the occurrence of investment mistakes – both delaying optimal investments as well as making inefficient investments.

6 Conclusion

This paper examines how the decision to invest in residential energy efficiency improvements is affected by subsidies for the upgrades as well as information (audits). While a number of studies explore households' decisions to participate in subsidized residential energy efficiency programs (e.g., Allcott and Greenstone (2017), Palmer and Walls (2015)), our analysis incorporates an important and unexplored dimension. Rather than focusing solely on whether or not a household participates in an energy efficiency program, we seek to understand the timing of households' participation decision – i.e. do households elect to receive in-home energy audits and make subsequent energy efficiency improvements early in their tenure in a home or after they have lived in the home for quite some time?

Ultimately, our analysis highlights that the timing margin – that is, when households make upgrades – is an important margin on which to focus. This is due to the fact that households that have just moved into a home and households that have lived in their homes for many years have potentially very different information sets. Households that have just moved into a home do not know how much energy they will consume in their home or what the resulting comfort level will be. Consequently, they will have little certainty surrounding the returns to making costly energy efficiency investments. In contrast, households that have lived in their present home for some time will have much more precise beliefs regarding the benefits of performing energy efficiency upgrades.

To shed light on how this inherent uncertainty surrounding the benefits affects households' decisions to participate in energy efficiency programs, we introduce a theoretical model of a household's decision to receive an energy audit and invest in subsequent energy efficiency upgrades. Importantly, the model captures that households may be uncertain about their baseline energy usage (and level of comfort) absent making investments in energy efficiency. In addition, we incorporate the fact that households face uncertainty in the share of energy use that could be reduced by performing upgrades. While a typical audit will remove much of the uncertainty surrounding the expected share of energy use that could be avoided by investing in energy efficiency, households will ultimately require time living in their homes – experiencing the level of comfort and observing their energy bills – to precisely understand the true benefits of performing upgrades.

By incorporating time varying uncertainty in the decision process, we provide new insights surrounding the efficacy of subsidizing energy efficiency upgrades. Importantly, we find that, while subsidies will induce more households to perform upgrades, much of this spending will be quite wasteful. In particular, a subset of the marginal participants will simply have their energy efficiency investments pulled-forward in time. That is, over time, as uncertainty surrounding the benefits of energy efficiency was resolved, these households would have elected to perform the upgrades without the additional subsidy. Even more troubling, we find that a portion of the marginal participants are induced to make investment mistakes – i.e. the costs of the upgrades they perform exceed the stream of social benefits they will provide.

While our results highlight that residential energy efficiency subsidies are a very costly way to encourage energy efficiency investments, the insights from our model do point to a potentially important margin policymakers can instead focus financial support towards. In particular, the findings from our analysis suggest that welfare gains could be achieved by reducing the uncertainty households face when deciding to make energy efficiency upgrades. Moreover, our results suggest that it would be particularly valuable to provide households with more precise information at the point when they move into a house. Households face substantial uncertainty in expected energy costs when moving into a new house because they don't yet know how their energy usage behaviors interact with the durable goods stock. Redesigning audit programs to predict a household's average use in a particular dwelling in addition to the percentage savings from specific upgrades would greatly reduce investment mistakes. This would be equivalent to receiving a large number of signals all at once rather than waiting to receive one noisy signal in each billing cycle. Such a program is also feasible given advances in machine learning. This recent-mover margin appears to be a particularly under-studied and under-exploited margin on which policymakers can focus. Not only is this the point in time when uncertainty is the greatest, our results also suggest it is when the likelihood of performing upgrades peaks.

References

Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market." *American Economic Review*, 105(8): 2501–38.

Allcott, Hunt, and Michael Greenstone. 2017. "Measuring the welfare effects of residential energy efficiency programs." National Bureau of Economic Research.

Allcott, Hunt, and Nathan Wozny. 2014. "Gasoline prices, fuel economy, and the energy paradox." *Review of Economics and Statistics*, 96(5): 779–795.

Boomhower, Judson, and Lucas W Davis. 2014. "A credible approach for measuring inframarginal participation in energy efficiency programs." *Journal of Public Economics*, 113: 67–79.

Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram. 2017. "Machine learning from schools about energy efficiency." National Bureau of Economic Research.

Busse, Meghan R, Christopher R Knittel, and Florian Zettelmeyer. 2013. "Are consumers myopic? Evidence from new and used car purchases." *American Economic Review*, 103(1): 220–56.

Chandra, Ambarish, Sumeet Gulati, and Milind Kandlikar. 2010. "Green drivers or free riders? An analysis of tax rebates for hybrid vehicles." *Journal of Environmental Economics and Management*, 60(2): 78–93.

Davis, Lucas W, and Gilbert E Metcalf. 2016. "Does better information lead to better choices? Evidence from energy-efficiency labels." *Journal of the Association of Environmental and Resource Economists*, 3(3): 589–625.

Dixit, Avinash K, Robert K Dixit, and Robert S Pindyck. 1994. *Investment under uncertainty.* Princeton university press.

Dubin, Jeffrey A, and Daniel L McFadden. 1984. "An econometric analysis of residential electric appliance holdings and consumption." *Econometrica*, 345–362.

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." *American Economic Review*, 105(5): 201–204.

Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram. 2018. "Do energy efficiency investments deliver? Evidence from the weatherization assistance program." *The Quarterly Journal of Economics*, 133(3): 1597–1644.

Gillingham, Kenneth, and Karen Palmer. 2014. "Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence." *Review of Environmental Economics and Policy*, 8(1): 18–38.

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

Holladay, Scott, Jacob LaRiviere, David Novgorodsky, and Michael Price. 2019. "Prices versus nudges: What matters for search versus purchase of energy investments?" *Journal of Public Economics*, 172: 151–173.

Houde, Sébastien. 2018. "How consumers respond to product certification and the value of energy information." *The RAND Journal of Economics*, 49(2): 453–477.

Hughes, Jonathan E, and Molly Podolefsky. 2015. "Getting green with solar subsidies: evidence from the California solar initiative." *Journal of the Association of Environmental and Resource Economists*, 2(2): 235–275.

Joskow, Paul L, and Donald B Marron. 1992. "What does a negawatt really cost? Evidence from utility conservation programs." *The Energy Journal*, 41–74.

Metcalf, Gilbert E, and Kevin 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.

Mian, Atif, and Amir Sufi. 2012. "The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program." *The Quarterly Journal of Economics*, 127(3): 1107–1142.

Newell, Richard G, and Juha Siikamäki. 2014. "Nudging energy efficiency behavior: The role of information labels." *Journal of the Association of Environmental and Resource Economists*, 1(4): 555–598.

Palmer, Karen, and Margaret Walls. 2015. "Limited attention and the residential energy efficiency gap." *American Economic Review*, 105(5): 192–95.

Table 1: Summary Statistics by Audit Uptake

	No-Audit ($N = 86,639$)		Audit ($N = 2,152$)			
	Mean	Std. Dev.	Mean	Std. Dev.	Diff.	p-value
Year Built	1971	31	1969	24	-2.07	0.002
Square Footage	1,736	882	1,912	956	176.8	0.000
Bedrooms	2.96	0.99	3.11	1.06	0.15	0.000
Floors	1.24	0.43	1.27	0.45	0.03	0.001
Mover (1 =yes)	0.08	0.27	0.14	0.35	0.06	0.000
Value ($'s)	148,101	134,145	170,272	143,937	22,173	0.000

Table 2: Immediate vs. Late Auditers

	Audit Within 30 Days ($N = 48$)		Audit After 30 Days ($N = 183$)			
	Mean	Std. Dev.	Mean	Std. Dev.	Diff.	p-value
Year Built	1964	26	1968	25	3.96	0.33
Square Footage	1,862	941	2,016	1,052	176.8	0.36
Bedrooms	3.69	4.38	3.15	0.98	-0.53	0.13
Floors	1.21	0.41	1.29	0.45	0.08	0.26
Value ($'s)	169,444	166,547	192,036	160,842	22,593	0.39

Table 3: Linear Probability Model of Audit Timing

	Moved in Sample				No Move in Sample			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Concurrent Month:								
Electricity$_t$ *(10 kWh/day)*	0.023***	0.026***	0.011	0.006	-0.004	-0.004	0.003	-0.004
	(0.008)	(0.010)	(0.011)	(0.012)	(0.006)	(0.005)	(0.007)	(0.007)
Gas$_t$ *(1 therm/day)*	0.007	0.013	-0.025**	-0.001	0.005	0.010	-0.011	0.007
	(0.010)	(0.016)	(0.012)	(0.013)	(0.006)	(0.008)	(0.007)	(0.010)
Previous Month:								
Electricity$_{t-1}$ *(10 kWh/day)*			0.032***	0.037**			-0.011*	0.0003
			(0.011)	(0.019)			(0.005)	(0.005)
Gas$_{t-1}$ *(1 therm/day)*			0.053***	0.030**			0.018**	0.005
			(0.014)	(0.014)			(0.008)	(0.011)
Premise FE	Y	Y	Y	Y	Y	Y	Y	Y
Month FE	N	Y	N	Y	N	Y	N	Y
N	698	698	698	698	14,507	14,507	14,507	14,507
R^2	0.01	0.06	0.05	0.08	0.001	0.02	0.01	0.08

Each model is estimated using monthly observations from households that elect to receive an audit more than one month after moving into a house. The dependent variable in each model is an indicator variable identifying the billing cycle in which a household elects to have an IHEA. The standard errors are robust to clustering at the household level and by month-of-sample. ** = Significant at the 5% level; *** = Significant at the 1% level.

Table 4: Likelihood of Audit Uptake by Year Built

	(1)	(2)	(3)	(4)
Year Built	-0.00002	-0.0002***	-0.0001***	-0.0001***
	(0.00004)	(0.00005)	(0.00004)	(0.00005)
Year Built×Mover	-0.001***	-0.002***	-0.001***	-0.0004**
	(0.0002)	(0.0003)	(0.0003)	(0.0002)
Premise Characteristics	N	Y	Y	Y
Zip Code FE	N	N	Y	Y
N	55,485	55,485	55,485	55,361
R^2	0.003	0.01	0.01	0.01

The linear probability models are estimated using each premise in our sample constructed between 1950 and 2000. Movers are defined as premises that are sold at any point during our three year sample. Premise characteristics include interactions between square footage bins, bedrooms, a multi-story indicator, and the mover indicator variable. The postal fixed effects include 35 zip code indicators. The standard errors are robust to heteroskedasticity. ** = Significant at the 5% level; *** = Significant at the 1% level.

Table 5: Likelihood of Installations by Year Built

	Movers		Non-Movers		All Homes	
	(1)	(2)	(3)	(4)	(5)	(6)
Year Built	-0.004*	-0.005**	-0.002**	-0.001	-0.002**	-0.001
	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Year Built×Mover					-0.002	-0.004
					(0.002)	(0.003)
Premise Characteristics	N	Y	N	Y	N	Y
Zip Code FE	N	Y	N	Y	N	Y
N	243	243	1,387	1,387	1,630	1,630
R^2	0.02	0.13	0.004	0.03	0.01	0.04

The linear probability models are estimated using each audited premise in our sample constructed between 1950 and 2000. Movers are defined as premises that are audited at any point after being sold during our three year sample. Premise characteristics include bins separating homes by square-footage, bedrooms, and a multi-story indicator variable. The postal fixed effects include 35 zip code indicators. The standard errors are robust to heteroskedasticity. * = Significant at the 10% level; ** = Significant at the 5% level; *** = Significant at the 1% level.

APPENDIX – For Online Publication

A Model Appendix

A.1 Derivation of equation (7)

The ex ante value function for $t+1$, $E[V_{t+1}]$, has two parts. First, over the range of support of $\mu(\theta)$ in which the household's posterior for the mean falls below the ex ante critical belief, $m_{t+1}(\theta) < \widetilde{m}$, the household will not audit and expects to receive the present value of wealth net of bills, conditional on the mean bill being below $\widetilde{m}$. Put another way, if expected energy bills aren't high, a household has less incentive to get an audit because install savings are expected to be lower.

The second part of the value function occurs over the range of the support of $\mu(\theta)$ in which the posterior exceeds the ex ante critical belief, $m_{t+1}(\theta) \geq \widetilde{m}$, the household will audit and receive their expected value given an audit. This second part also has two components: the support of α over which an install occurs, and the support of α over which it does not, as defined in equation (6). As a result, the ex ante expected value function is:

$$E[V_{t+1}] = \int_{\widetilde{m}}^{\infty} \Bigg[- A + \Big(1 - G(\bar{\alpha})\Big) \frac{w - p\mu}{1-\delta} + \int_0^{\bar{\alpha}} \Big(\frac{w - \alpha p \mu}{1-\delta} - F\Big) dG(\alpha) \Bigg] dH_{t+1}(\mu) + \int_{-\infty}^{\widetilde{m}} \frac{w - p\mu}{1-\delta} dH_{t+1}(\mu). \quad \text{(A1)}$$

By noting that the expected value of m_{t+1} at time t is m_t and evaluating the integrals, equation (A1) immediately simplifies to equation (7):

$$E[V_{t+1}] = \underbrace{\frac{w}{1-\delta}}_{\text{present value of wealth}} - \underbrace{A \cdot \Big(1 - H(\widetilde{m})\Big)}_{\text{ex ante expected audit cost}} - \underbrace{F \cdot G(\bar{\alpha})\Big(1 - H(\widetilde{m})\Big)}_{\text{ex ante expected install cost}} - \underbrace{\frac{p}{1-\delta}\Bigg(1 - G(\bar{\alpha}) \cdot (1 - \widehat{\alpha})\Bigg)(1 - H(\widetilde{m}))\widehat{\mu}_H}_{\text{expected bills if audit, given uncertainty over install}} - \underbrace{\frac{p}{1-\delta} H(\widetilde{m}) \widehat{\mu}_L}_{\text{expected bills if don't audit}} . \quad \text{(A2)}$$

A.2 Proof of Proposition 1

At $t = 0$, all households with initial prior $m(\theta) \geq \widetilde{m}$ request an audit. We can express this share of households as

$$\int_{\widetilde{m}}^{\infty} dM(m(\theta)) = 1 - M(\widetilde{m}),$$

with $M(\widetilde{m})$ the share of remaining un-audited households in time $t = 1$.

At $t = 1$, the probability that a particular household with initial prior $m(\theta) < \widetilde{m}$ updates their belief to $m_1 \geq \widetilde{m}$ is given by

$$Pr(m_1(\theta) \geq \widetilde{m}) = 1 - H_1(\widetilde{m}; \tau + 1 \cdot r, \theta)$$

where we note that the distribution of the posterior has an increasing precision and depends on characteristics θ. In order to obtain the share of total households that audit in $t = 1$, we integrate this expression over the distribution of home types with initial priors that were not high enough to justify an audit at $t = 0$:

$$N_1 = \int_{-\infty}^{\widetilde{m}} (1 - H_1(\widetilde{m}; \tau + 1 \cdot r, \theta))\, dM(m(\theta))$$

At $t = 2$, we can similarly calculate the share of households for a given home type that choose to audit based on their updated belief ($Pr(m_2(\theta) \geq \widetilde{m})$) and then aggregate over home types. However, we also need to adjust for the fact that, with probability $H_1(\widetilde{m}; \tau + 1r, \theta)$ a home of type θ did not audit at $t = 1$, so that only $(1 - N_1 - M(\widetilde{m}))$ of the households remain unaudited. This expression is given by

$$N_2 = \int_{-\infty}^{\widetilde{m}} \Big(1 - H_2(\widetilde{m}; \tau + 2r, \theta)\Big) \cdot H_1(\widetilde{m}; \tau + 1r, \theta) dM(m(\theta))$$

where we have adjusted the share of total homes auditing at time t by the probability that they did not audit in the previous period. We can now write the general expression for the share of total households that choose to audit at a given time t:

$$N_t = \int_{-\infty}^{\widetilde{m}} \Big(1 - H_t(\widetilde{m}; \tau + tr, \theta)\Big) \prod_{s=1}^{t-1} H_s(\widetilde{m}; \tau + sr, \theta) dM(m(\theta)) \tag{A3}$$

Equation (A3) is declining in t for two reasons. First, the term $\prod_{s=1}^{t-1} H_s(\widetilde{m}; \tau + sr, \theta)$ is clearly declining in time because $H_s < 1$ for any s. Second, the precision of the posterior distribution increases over time such that H_t is a mean-preserving spread of H_{t+1}, which implies that, for households that have not audited as of time t, $(1 - H_t) > (1 - H_{t+1})$.

A.3 Proof of Lemma 2

By assumption, the true mean is increasing in θ: $m'(\theta) > 0$. From the expression for updated priors $m_t(\theta) = \frac{\tau m(\theta)+tr\bar{e}}{\tau+tr} = \frac{\tau m(\theta)+r\sum_{s=0}^{t} e_s}{\tau+tr}$

$$\frac{\partial m_t(\theta)}{\partial \theta} = \frac{\tau m'(\theta)}{\tau + tr} > 0$$

$$\frac{\partial m_t(\theta)}{\partial e_t} = \frac{r}{\tau + tr} > 0$$

Both expressions are clearly declining in t.

A.4 Derivation of equation (9)

We first simplify (8) and then differentiate it with respect to θ (which produces an almost identical result to differentiating with respect to e_t). We then simplify the expression to derive equation (9).

Simplifying (8) slightly:

$$\begin{aligned} E[V_t(audit)] - E[V_t(delay)] &= G(\bar{\alpha})(1-\widehat{\alpha}) \cdot pm_t \\ &\frac{\delta}{1-\delta} G(\bar{\alpha})(1-\widehat{\alpha})H(\widetilde{m})p\widehat{\mu}_L - \Big(A + F \cdot G(\bar{\alpha})\Big) \cdot \left(1 - \delta\Big(1 - H(\widetilde{m})\Big)\right) \end{aligned} \quad \text{(A4)}$$

Differentiating (A4) with respect to θ gives

$$\begin{aligned} \frac{\partial E[V_t(audit)] - E[V_t(delay)]}{\partial \theta} &= G(\bar{\alpha})(1-\widehat{\alpha})pm'_{t,\theta} + pm_t \frac{\partial G(\bar{\alpha})(1-\widehat{\alpha})}{\partial m_t} m'_{t,\theta} \\ &+ \frac{\delta p}{1-\delta}\left(G(\bar{\alpha})(1-\widehat{\alpha}) \frac{\partial H(\widetilde{m})\widehat{\mu}_L}{\partial m_t} + \frac{\partial G(\bar{\alpha})(1-\widehat{\alpha})}{\partial m_t} H(\widetilde{m})\widehat{\mu}_L \right) m'_{t,\theta} \\ &+ \delta\Big(A + F \cdot G(\bar{\alpha})\Big) \frac{\partial (1 - H(\widetilde{m}))}{\partial m_t} m'_{t,\theta} - F\Big(1 - \delta(1 - H(\widetilde{m}))\Big) \frac{\partial G(\bar{\alpha})}{\partial m_t} m'_{t,\theta} \end{aligned} \quad \text{(A5)}$$

Note that an increase in the prior shifts the distributions of μ and α. We therefore need to derive

$$\frac{\partial (1 - H(\widetilde{m}))\widehat{\mu}_H}{\partial m_t}, \quad \frac{\partial H(\widetilde{m})\widehat{\mu}_L}{\partial m_t}, \quad \frac{\partial (1 - H(\widetilde{m}))}{\partial m_t}, \quad \frac{\partial G(\bar{\alpha})}{\partial m_t}, \quad \text{and } \frac{\partial G(\bar{\alpha}) \cdot (1-\widehat{\alpha})}{\partial m_t}$$

We first note that $h(\mu)$ is the Normal pdf with standard deviation $\sigma = 1/(\tau + tr)$, and use integration by parts to derive

$$\begin{aligned}
\frac{\partial(1-H(\widetilde{m}))\widehat{\mu}_H}{\partial m_t} &= \frac{\partial}{\partial m_t}\int_{\widetilde{m}}^{\infty}\mu h(\mu)d\mu \\
&= \int_{\widetilde{m}}^{\infty}\mu\frac{\partial h(\mu)}{\partial m_t}d\mu \\
&= \int_{\widetilde{m}}^{\infty}\frac{\mu}{\sigma}\frac{\mu-m_t}{\sigma}h(\mu)d\mu \\
&= \frac{1}{\sqrt{2\pi}\sigma}\int_{\widetilde{m}}^{\infty}\Big(z+\frac{m_t}{\sigma}\Big)z\exp(-z^2/2)dz, \quad \text{where } z=\frac{\mu-m_t}{\sigma} \\
&= \frac{1}{\sqrt{2\pi}\sigma}\int_{\widetilde{m}}^{\infty}udv, \quad \text{where } (z+\frac{m_t}{\sigma}), dv = z\exp(-z^2/2)dz \\
&= \frac{1}{\sqrt{2\pi}\sigma}\Big[-(z+\frac{m_t}{\sigma})\exp(-z^2/2)\Big]_{\widetilde{m}}^{\infty} - \frac{1}{\sqrt{2\pi}\sigma}\int_{\widetilde{m}}^{\infty}-\exp(-z^2/2)dz \\
&= \widetilde{m}\frac{h(\widetilde{m})}{\sigma}+(1-H(\widetilde{m})) \\
&= \widetilde{m}h(\widetilde{m})(\tau+tr)+(1-H(\widetilde{m}))
\end{aligned}$$

A similar procedure shows that

$$\frac{\partial H(\widetilde{m})\widehat{\mu}_L}{\partial m_t} = -\widetilde{m}h(\widetilde{m})(\tau+tr)+H(\widetilde{m})$$

And

$$\frac{\partial(1-H(\widetilde{m}))}{\partial m_t} = h(\widetilde{m})(\tau+tr)$$

Recalling that $\bar{\alpha}=\bar{\alpha}(m_t)$ with $\bar{\alpha}'(m_t)>0$, it is clear that

$$\frac{\partial G(\bar{\alpha})}{\partial m_t} = g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}$$

We can then use the Leibniz rule to derive

$$\begin{aligned}
\frac{\partial G(\bar{\alpha})\cdot(1-\widehat{\alpha})}{\partial m_t} &= \frac{\partial}{\partial m_t}\Big[G(\bar{\alpha})-\int_0^{\bar{\alpha}(m_t)}\alpha dG(\alpha)\Big] \\
&= g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}-\bar{\alpha}g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t} \\
&= (1-\bar{\alpha})g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}
\end{aligned}$$

Plugging these expressions into equation (A5) yields

$$\begin{aligned}
\frac{\partial E[V_t(audit)]-E[V_t(delay)]}{\partial\theta} &= G(\bar{\alpha})(1-\widehat{\alpha})pm'_{t,\theta}+pm_t(1-\bar{\alpha})g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}m'_{t,\theta} \\
&+\frac{\delta p}{1-\delta}\Bigg(G(\bar{\alpha})(1-\widehat{\alpha})\Big(-\widetilde{m}h(\widetilde{m})(\tau+tr)+H(\widetilde{m})\Big)+(1-\bar{\alpha})g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}H(\widetilde{m})\widehat{\mu}_L\Bigg)m'_{t,\theta} \\
&+\delta\Big(A+F\cdot G(\bar{\alpha})\Big)h(\widetilde{m})(\tau+tr)m'_{t,\theta}-F\Big(1-\delta(1-H(\widetilde{m}))\Big)g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}m'_{t,\theta} \qquad \text{(A6)}
\end{aligned}$$

We can simplify this further by noting that $F=\frac{(1-\bar{\alpha})pm_t}{1-\delta}$ and $m_t=H(\widetilde{m})\widehat{\mu}_L+(1-H(\widetilde{m}))\widehat{\mu}_H$ and combining like terms to find

$$\frac{\partial E[V_t(audit)] - E[V_t(delay)]}{\partial\theta} = \left[\frac{p}{1-\delta}G(\bar{\alpha})(1-\widehat{\alpha})\Big(1-\delta+\delta H(\widetilde{m})\Big)+\delta h(\widetilde{m})(\tau+tr)A\right]m'_{t,\theta}$$
$$+\frac{\delta p}{1-\delta}\left[h(\widetilde{m})(\tau+tr)G(\bar{\alpha})\Big((1-\bar{\alpha})m_t-(1-\widehat{\alpha})\widetilde{m}\Big)+g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}(1-\bar{\alpha})(1-H(\widetilde{m}))(m_t-\widehat{\mu}_H)\right]m'_{t,\theta} \tag{A7}$$

A.5 Proof of Proposition 4

We have established that the household makes an installation following an audit if the present value of installation benefits, denoted b_I, exceeds the installation cost:

$$b_I = \frac{1-\alpha}{1-\delta}pm_t(\theta) \geq F_t$$

For larger θ or e_t we have

(i) $\frac{\partial b_I}{\partial\theta} = \frac{1-\alpha}{1-\delta}pm'_{t,\theta}$

(ii) $\frac{\partial b_I}{\partial e_t} = \frac{1-\alpha}{1-\delta}pm'_{t,e_t}$

From Lemma 2, $m'_{t,\theta}$ and m'_{t,e_t} are positive and declining in t.

A.6 Impact of Audit Subsidy on Additionality and Mistakes

By applying Jensen's Inequality to the ex post and ex ante usage threshold, we can see that an audit itself lowers the usage threshold for a household to make an installation, even in the absence of an installation subsidy. In other words, the very fact of an audit makes a household more likely to make an install even if the audit teaches the household nothing about their true mean use. To see this, recall that the ex post usage threshold at the average level of savings $E(\alpha)$ is a convex function of $E(\alpha)$:

$$\bar{m}(E(\alpha)) = \frac{F}{p}\frac{1-\delta}{1-E(\alpha)}$$

,

In the absence of an audit program, the household has no way of learning about the savings rate α from an install, and will make an energy efficient installation if their prior exceeds the ex ante usage threshold:

$$E(\bar{m}) = \int_0^1 \frac{F}{p}\frac{1-\delta}{1-\alpha}dG(\alpha)$$

.

By Jensen's Inequality, the ex post usage threshold at the average level of savings α is lower than the ex ante usage threshold which a household would apply to their installation decision without an audit ($\bar{m}(E(\alpha)) < E(\bar{m})$). In other words, an audit induces installation for households whose prior estimate of mean use is between $\bar{m}(E(\alpha))$ and $E(\bar{m})$.

B Energy Consumption Variability

A key policy relevant aspect of our analysis is how predictable a home's electricity usage is in any given month. If homes with similar physical characteristics have wildly different observed electricity consumption, then some homeowners could move into homes that appear to be "energy hogs" based upon observables and make an install based upon expected savings only to discover the home was actually energy efficient and the install was a mistake. Put more bluntly, as the variance in home energy usage conditional on observable increases, the value of gaining information about the actual energy usage of the home increases. Of course preferences also matter since some households are simply willing to pay more for services provided by electricity use; we return to this point shortly.

To investigate how predictable a home's electricity usage is we perform a simple statistical exercise. For the sample of 88,791 homes in the data we trim the sample to only include the homes between the 10th and 90th percentile of square footage. We then run a regression of monthly electricity usage on linear and squared values of year built, number of bedrooms, a 3rd degree polynomial of square footage, number of floors and indicator variables for month and year. The relatively simple regression is meant to mirror what a sophisticated home buyer might use to form expectations of average electricity usage of a home. From that regression we get predicted values. Next we take the median year built, median number of bedrooms, median number of floors and homes 5% above and 5% below the median level of square feet in the sample. This leaves 236 unique homes.

Figures A1 and A2 show observed and predicted electricity usage of the 238 households closest to the median household in the data for two shoulder months, April and October, in 2012. We pick these months because they are mild and thus should have the least variation in usage in absolute terms. Observed electricity usage varies wildly, from less than 200 kWh up to 2,000 kWh, with a mean of roughly 1,000 kWh. Conversely, predicted values range from about 800 kWh to almost 900 kWh. The point of this exercise is to show that household electricity usage varies a great deal. Even if 50% of the observed difference in electricity consumptions is due to preferences, the remaining range of observed electricity differences would still be large relative to the mean. We take this as evidence that individual homes have high variation in their energy efficiency even conditional on observables.

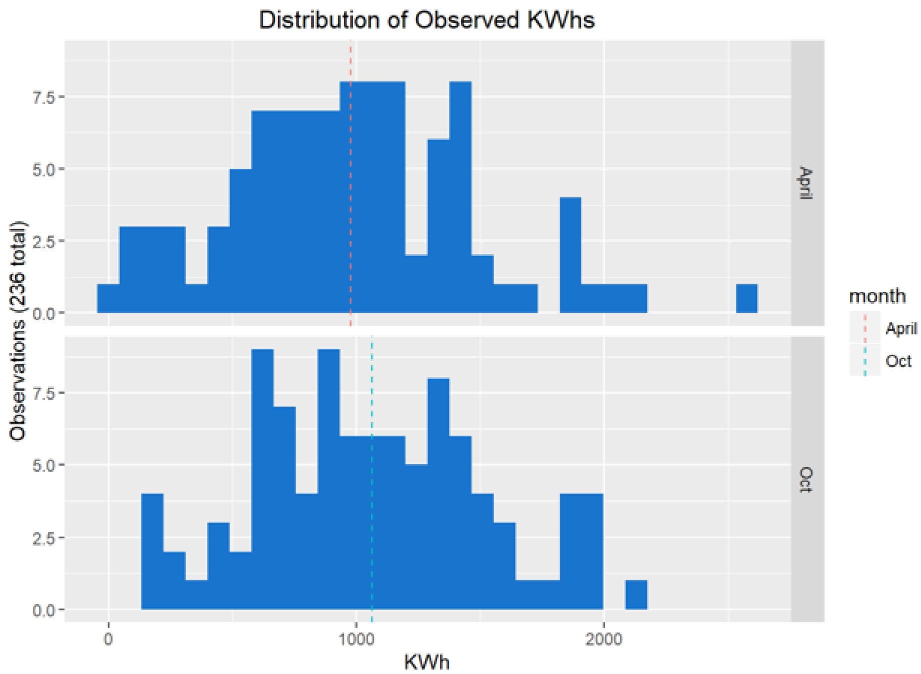

Figure A1: Notes: Observed April and October 2012 electricity usage.

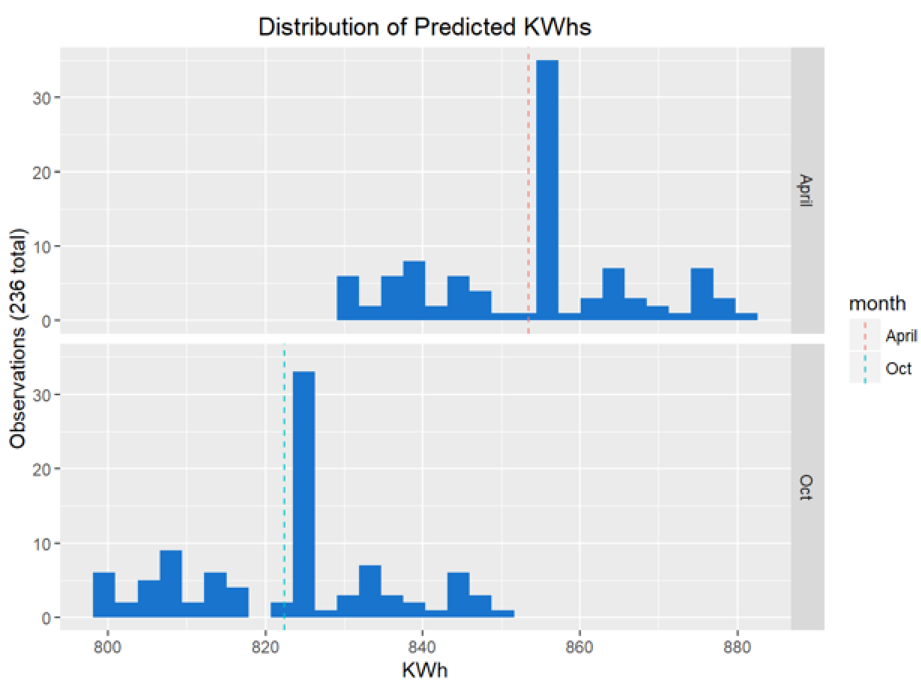

Figure A2: Notes: Predicted April and October 2012 electricity usage.

C Appendix Tables and Figures

Table A1: Frequency of Upgrades Conditional on Audit

	Movers ($N = 308$)	Non-Movers ($N = 1,844$)		
	Mean	Mean	Diff.	p-value
ANY Upgrades	0.63	0.62	0.01	0.64
Primary Windows	0.41	0.40	0.002	0.94
HVAC Replacement	0.06	0.05	0.01	0.57
HVAC Tune-up	0.02	0.02	0.000	0.96
Duct Repair/Replace	0.02	0.04	-0.02	0.17
Duct Sealing	0.12	0.11	0.01	0.66
Attic Insulation	0.12	0.14	-0.02	0.27
Air Sealing	0.10	0.11	-0.02	0.38
Wall Insulation	0.03	0.03	-0.01	0.62
Floor/Perimeter Insulation	0.04	0.04	-0.002	0.89
Vapor Barrier	0.02	0.02	-0.002	0.85
General Rehab.	0.01	0.01	-0.01	0.25
Water Heater Insulation	0.003	0.01	-0.003	0.49
Water Pipe Insulation	0.003	0.01	-0.01	0.29

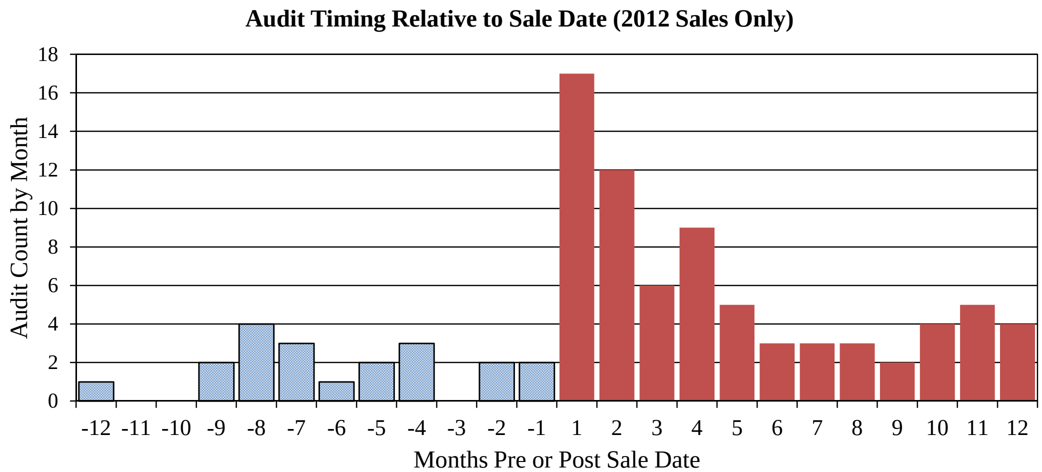

Figure A3: Notes

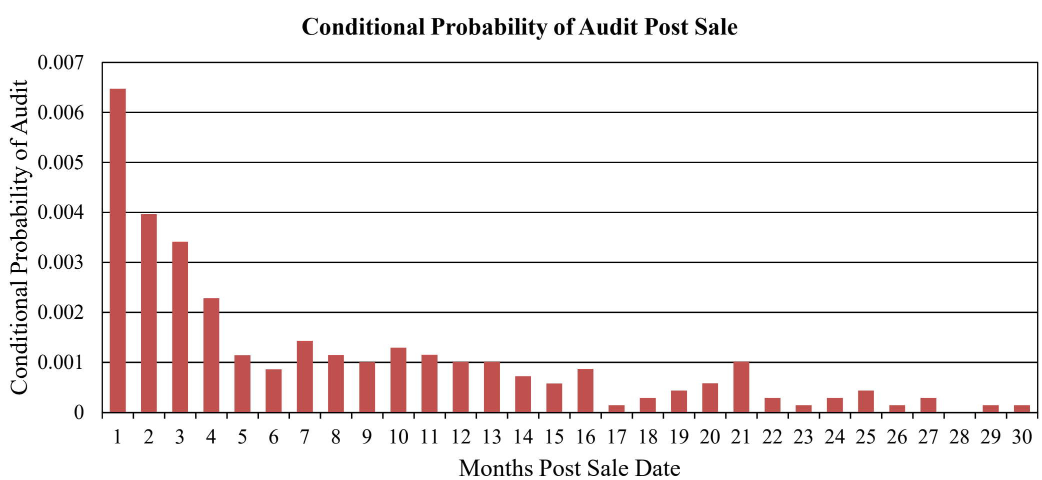

Figure A4: Notes

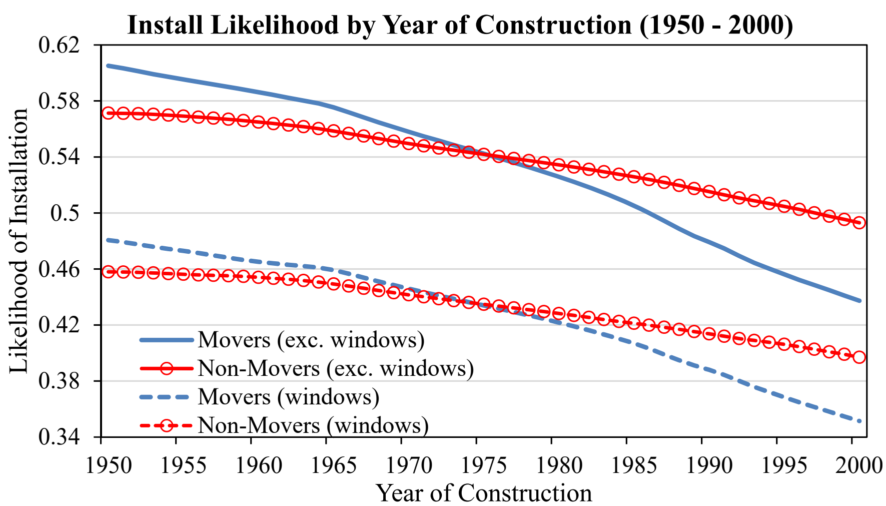

Figure A5: Notes