Uncertainty and Additionality in Energy Efficiency Programs*

Ben Gilbert,[†] Jacob LaRiviere,[‡]and Kevin Novan[§]

October 22, 2020

Abstract

Programs subsidizing residential energy efficiency investments often suffer from poor additionality. In this paper, we demonstrate that an additional inefficiency can arise due to the fact that there is uncertainty surrounding the benefits achieved by performing the upgrades. Introducing a dynamic model of a household's decision to invest in energy efficiency upgrades, we highlight that, when faced with uncertainty, a subset of households will make irreversible investments 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 upgrades, 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.

JEL Codes: 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. Catie Hausman, Alan Krupnick, Karen Palmer, and Michael Price made comments which significantly improved this paper. We also thank Brittany Tarufelli for excellent research assistance.

†Division of Economics & Business at Colorado School of Mines. E-mail: *bgilbert@mines.edu*.

‡Microsoft, Department of Economics at University of Tennessee, & Department of Economics at University of Washington. E-mail: *jlariv@microsoft.com*.

§Department of Agricultural and Resource Economics at UC Davis. E-mail: *knovan@ucdavis.edu*.

1 Introduction

To reduce the fuel, capital, and environmental costs required to satisfy residential energy demand, a substantial amount of money is being poured into programs subsidizing energy efficiency.[1] A common concern with these programs is that they suffer from notoriously poor additionality. That is, many of the subsidized upgrades would have occurred without financial support, suggesting that public funds are being used very inefficiently (Joskow and Marron (1992), Boomhower and Davis (2014), Globus-Harris (2019)).[2]

Thus far, the literature studying additionality has effectively viewed the issue as a classic principal agent problem with a single source of uncertainty. Specifically, governments subsidizing energy efficiency upgrades are uncertain about homeowners' demand for the upgrades. In practice, however, there is another potentially important source of uncertainty – households themselves don't know how much energy will be saved, or how much comfort will be increased, when investing in upgrades. In this paper, we introduce a dynamic model of a household's investment decision that incorporates uncertainty in the benefits achieved by making energy efficiency upgrades. By doing so, we are able to provide a more nuanced understanding of additionality. In particular, we highlight that, not only are many participants in energy efficiency programs inframarginal, or 'non-additional', many of the 'additional' participants are encouraged to make investments with upfront costs that exceed the stream of future benefits. This insight suggests that energy efficiency programs are even less efficient than previously thought.

Our analysis contributes to the literature examining the decision to participate in energy efficiency programs (e.g., Holladay et al. (2019), Allcott and Greenstone (2017),

[1] Much of this support is funneled through customer incentive programs and subsidies offered by energy utilities. From 2013 through 2017, U.S. electric utilities spent \$7.5 billion on residential energy efficiency programs. 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. Similar types of financial support also come from the federal government. 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/`.

[2] Low additionality has also been documented 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)), appliance subsidies (Houde and Aldy (2017)), as well as subsidies for hybrid vehicles (Chandra, Gulati and Kandlikar (2010)).

Palmer and Walls (2015)). While previous studies focus exclusively on the choice of whether or not to participate in an energy efficiency program, we pay particular attention to the timing of the participation decision. First and foremost, we document that the timing of participation in residential energy efficiency programs is not random. Using information from an energy efficiency program in a medium-sized MSA in the U.S., we compare the date that homes are sold to the date that households elect to participate in the program. We observe that the program participation rate is more than 10 times higher during the weeks immediately after a move compared to the baseline participation rate.

Ultimately, our goal with this analysis is not to explain why participation in energy efficiency programs peaks right after moving into a home.[3] Rather, we are interested in understanding how the timing of participation affects the economic efficiency of the investments that are being made. Intuitively, the timing may be quite important due to the fact that households 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 energy they will consume as well as what the resulting comfort level will be. Consequently, new homeowners will have the least certainty surrounding the benefits that would be achieved by making energy efficiency upgrades.

We use a dynamic model of the household's investment decision to explore how uncertainty affects a homeowner's decision to participate in an energy efficiency program and perform upgrades. In the model, households form priors about their future energy usage based on the observable characteristics of the homes (e.g., the age and condition of the home) and then update these priors with each new energy bill. Households that expect to consume more energy in the future expect to receive a larger stream of benefits from investing in energy efficiency upgrades. Consistent with typical energy efficiency programs offered by energy utilities, we model the participation and investment decision as a two-step process. In the first stage, a household must decide whether to receive an in-home energy audit. Through an audit, homeowners receive expert advice regarding what types

[3]There are of numerous factors that could cause program participation to spike immediately after moving into a home. For example, the non-pecuniary, fixed cost of performing upgrades, which previous studies find to be a sizable portion of the fixed cost of participating (e.g. Fowlie, Greenstone and Wolfram (2015) and Allcott and Greenstone (2017)), may be lower prior to fully moving into a home. In some cases, service providers also advertise the programs to knew customers.

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.

Unlike previous theories around investing in energy efficient durable goods which revolve around stochastic energy bills, our model adds uncertainty which resolves over time. The canonical rational agent investment models of Hassett and Metcalf (1992) and Hassett and Metcalf (1995) study households' decision to invest in energy efficient durable goods with stochastic electricity prices steadily rising around a long run trend. Like Hassett and Metcalf (1995), we have persistent, stochastic electricity bills but also include uncertainty over the expected future savings from making an energy efficiency investment which resolves over time. In particular, people moving into a home don't have a precise understanding of their typical energy use, and thus energy bills, in that house until they've been there for a while; they learn about their bills and potential gains from an energy efficiency investment over time.

Adding learning about a home's electricity usage over time allows us to provide two new insights surrounding the impacts of subsidizing energy efficiency upgrades and audits. First, we point out that subsidies pull forward some investments that would have occurred without financial support as uncertainty resolves in the future. While previous work empirically highlights that similar subsidies can lead to demand being pulled forward in time (Mian and Sufi (2012), Hughes and Podolefsky (2015) and Houde and Aldy (2017)), our model allows us to highlight the role that uncertainty can play in driving this outcome. Second, we demonstrate that, in the presence of unresolved uncertainty, subsidies will incentivize a subset of households to invest in upgrades with costs that exceed the future stream of social benefits they will provide.

The findings from our theoretical analysis suggest that the social benefits provided by subsidizing energy efficiency upgrades and audits may be far 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 suggests that many of these additional investments may not be truly additional, but rather simply pulled forward in time. Moreover, many of the additional participants may be overinvesting in energy efficiency because of reducible uncertainty about the future stream of benefits. We estimate the energy use uncertainty that new occupants of a home are likely to face, and use these estimates to parameterize our model. We find that more than one third of households that made upgrades with this level of uncertainty would realize negative net benefits ex post.

Our analysis also contributes to the literature examining how information affects consumers' investments in energy-related durables.[4] When purchasing an energy-consuming durable good (e.g., a vehicle, an appliance, a house), consumers face a trade-off between the upfront purchase price and the future stream of energy payments required to operate the good. Recognizing that inefficient investment decisions will occur without complete information regarding these future operating costs, policymakers have long focused on mandating the provision of information.[5] Recent work demonstrates that the provision of this information (Houde (2018), Allcott and Taubinsky (2015), and Newell and Siikamäki (2014)), as well as the precision of the information (Davis and Metcalf (2016)), can meaningfully alter consumers' appliance purchase decisions. Similarly, our analysis suggests that meaningful welfare gains could be achieved by focusing resources towards reducing the uncertainty surrounding the benefits from improving a home's energy efficiency. In particular, this information would be the most valuable when households are moving into a home – when the likelihood of performing inefficient energy efficiency upgrades peaks.

In the following section, we first provide evidence that the likelihood of participating in energy efficiency programs peaks immediately after moving into a home – precisely when we would expect the greatest uncertainty surrounding the benefits from investing in energy efficiency. Section three introduces the theoretical model used to examine household-level investment decisions in the face of uncertainty. Section four discusses the insights our

[4]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.

[5]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.

theoretical model provides with regards to subsidizing energy efficiency upgrades and home energy audits. Section five concludes.

2 Timing of Program Participation

Before exploring how the timing of participation in energy efficiency programs can affect the economic efficiency of the resulting investments, we first document that the relationship between the duration a household has lived in a home and the likelihood of participating in a residential energy efficiency program is not random. We focus specifically on a program run by an US energy utility in a medium sized MSA. From 2011 through 2013, the local electric and gas utility provided residential customers with the opportunity to pay a subsidized rate of $50 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.[6] Households choosing to make energy efficiency upgrades received up to $500 from local power providers in rebates. To receive the installation rebate, a household must have first had an audit.

To explore the timing of the participation decisions, we combine two key pieces of information. First, we observe the date that an audit was requested, as well as the date of any subsequent upgrades, at each premise.[7] Second, from assessor data, we observe each date that a premise is sold. In total, there are 150,658 unique premises in the utility service area. 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.[8] This leaves us with a final sample of 88,791, of which 2,573 elect to receive

[6] In the setting we study, there were no binding audit supply constraints. The data we observe records both the date the audit was requested and the date the audit occurred. In every case in which an audit was requested, an audit occurred within the same week, and typically within one day.

[7] The lag between the date that an audit was requested and date that it subsequently occurred was very short – uniformly under one week. The data was shared under a privacy agreement directly by the auditing agency.

[8] 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 err on the conservative side and drop premises

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

audits during the sample period.

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 mover indicator is equal to one if a premise is ever sold during the 2011-2013 period. 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 sample period.[9] Overall just under 15% of audits are scheduled by movers.

Focusing on the premises that were sold and audited at some point during our 3-year sample period, we calculate the number of days between the recorded sales date and the date the audit was scheduled. The top panel of Figure 1 displays the histogram of the number of months between the audit date and the sale date. There is a clear spike up in audits occurring during the first month following the home sale. Similarly, the bottom panel of Figure 1 displays the probability a home is audited during a given month after moving in, conditional on not being audited in the prior months. Again, the figure displays a dramatic spike in the probability of being audited immediately after moving into a home. Taken together, Table 1 and Figure 1 show that despite having observed zero or only a

which do not have a clear address match across the two samples.

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

few electricity or heating bills, households are more likely to get audits immediately after moving into a home.

We are also able to examine how often the audited households make subsequent upgrades. Table 2 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 our sample. Importantly, Table 2 highlights that energy efficiency upgrades were performed in 63% of the 308 homes that were sold and audited as well as in 62% of the audited homes that were not sold during our three year sample. Observing audited households electing not to perform any subsequent upgrades provides strong evidence that the information households receive through the audit is in fact being used. If households instead believed that the audit provided no reliable information, then only households that were determined to make efficiency upgrades would pay for the audit.

Table 2: 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

Table 2 also highlights that the vast majority of upgrades are intended to improve the thermal insulation of a home (e.g., new windows and attic insulation) or improve

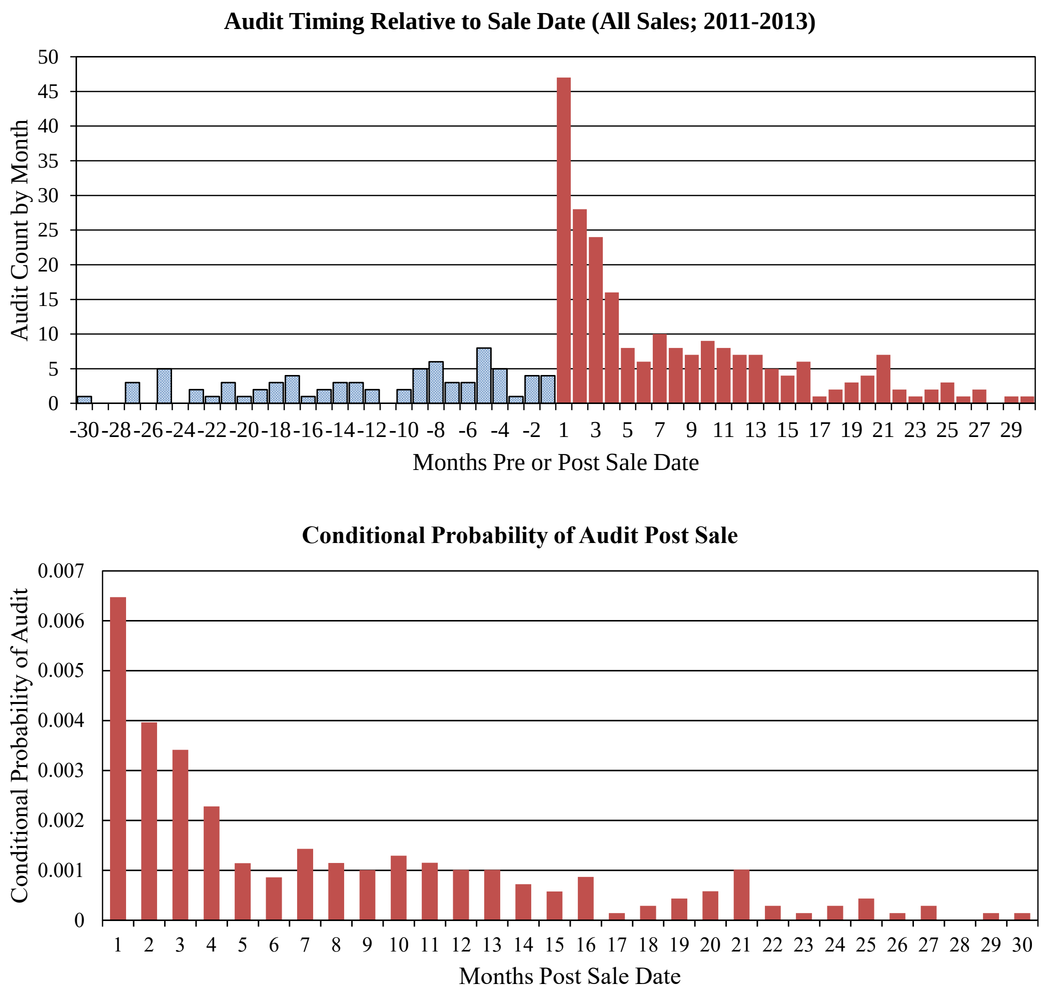

Figure 1: The top panel displays the histogram of number of Audits relative to move date for homes where an audit occurs and the home is sold between 2011 and 2013. The bottom panel displays the probability (frequency) of an audit occuring in a given month post-sale, conditional on no audit being observed in the preceding months.

the efficiency of the heating and cooling systems (e.g., duct sealing). These upgrades can provide private benefits to the homeowners in the form of reduced expenditures on heating and cooling – which account for over 50% of residential energy use in the U.S. – as well as increased comfort (e.g., less drafty homes).[10]

Ultimately, these private benefits are very difficult to predict prior to making an upgrade. Accurately estimating the energy savings that could be achieved by investing in insulation and heating/cooling upgrades requires answers to two critical questions. First, how much will the thermal insulation and heating and cooling efficiency be improved? Second, how intensively will the occupants of the home use air conditioning and heating? While engineering models used in the audit process can produce estimates of the thermal efficiency gains, any resulting estimates of the potential energy savings will be conditional on some assumed pattern of household behavior (e.g., the frequency of AC usage). Thus, an engineering model alone cannot perfectly predict future energy savings.

Instead, to understand the true magnitude of the benefits from making energy efficiency upgrades, households must also have a clear understanding of how much heating and cooling they will actually demand going forward. Intuitively, after a household has resided in a premise for a long period of time, they will likely have a great deal of certainty surrounding their energy use and comfort, and therefore, will likely have a fairly precise idea of the benefits that could be achieved by improving the efficiency of their heating and cooling system. In contrast, when a household is first moving into a home, they may have a poor understanding of how warm the home may get during the summer and how cold it may be during the winter. While they will likely form expectations based on the observed characteristics of the home (e.g., its age or it physical condition), the uncertainty surrounding future energy consumption, as well as the comfort, will likely be the greatest when a household is first moving into a home. Unfortunately, predicting any particular home's electricity use is challenging. In Appendix C we show that observed variation in the mean of monthly energy usage, conditional on observable housing stock characteristics, is very large (coefficient of variation of $\tilde{0}.45$). Even when we restrict the sample to the

[10] If the reduction in the cost of cooling and heating leads to a rebound in the consumption of cooling and heating services, then this would lead to increased comfort. To the extent that the upgrades are capitalized into the value of the home, there are also private benefits stemming from the increased resale value.

shoulder months when variability should be lower, the variation in usage of the median household in our dataset is still quite large (coefficient of variation of $\tilde{0}.3$). However, model performance improves materially when a household's lagged usage variables are included in the model.

In sum, households appear to be far more likely to participate in residential energy efficiency programs immediately after moving into a home. There are certainly a wide range of factors that could cause this pattern to emerge.[11] However, regardless of why this pattern emerges, it suggests that the upgrades performed through a residential energy efficiency program are more likely to occur when households face the greatest uncertainty surrounding the magnitude of the benefits. In the subsequent section, we present a dynamic model that allows us to explore how uncertainty surrounding the benefits of investing in energy efficiency upgrades affects households' decisions to participate in energy efficiency programs and perform subsequent upgrades.

3 Theoretical Model

3.1 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 upgrade. The only way to learn about upgrade savings is by performing an audit. Expected benefits of an upgrade are the product of mean energy bills and savings from an upgrade.

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

[11] New owners likely have different preferences that the previous occupants. As a result, upon moving in, the new owners are likely to make a wave of changes and investments (e.g., remodeling, changing appliances, upgrading the energy efficiency). Another potential factor is that there are likely sizable 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 an inspection and upgrades performed on the home. While it was not a factor in the specific program being studied, in some cases there will be informational campaigns/marketing targeted at new households moving into a utility's service area.

(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 larger homes, ceteris paribus). To summarize, a household moving into a type θ dwelling forms a prior estimate, $m(\theta)$, about mean energy use which is an unbiased estimate of their true mean,

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

$\mu(\theta)$, 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:

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

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 $H_t(\mu)$.[13]

Allowing households to update priors about their home's mean electricity use is the key distinction between our model and the earlier work of Hassett and Metcalf (1992) and Hassett and Metcalf (1995). Hassett and Metcalf (1995) assumes electricity prices are rising stochastically around a trend. The decision to make an efficiency upgrade is based upon the stochastic process of electricity prices, or electricity bills in our case. In our model we also have reducible parametric uncertainty over parameters dictating a home's stochastic electricity bills (e.g., $m_t(\theta)$ is a prior for $\mu(\theta)$) through passive learning (LaRiviere et al. (2018)). This distinction allows us to highlight the case of moving in when a household has no information about the house other than observables.

We model energy efficiency upgrades as being available at a fixed cost κ, which reduce energy bills by $(1-\alpha)\%$ for some $\alpha \in (0,1)$. After making an upgrade, 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, we suppose that households can schedule an audit at any time for a fixed cost of A. We assume that households have

[13] In this paper, we don't directly observe households' priors, nor the speed they update beliefs. As a result, we can't measure the magnitude of this uncertainty relative to average electricity bills. We present empirical evidence in the Appendix that variation in electricity use is large even controlling for home characteristics. Hence uncertainty over a home's electricity profile is likely to be large at time of move in.

beliefs about α according to the distribution $G(\alpha)$. We assume $G(\alpha)$ is shared across all households.

We model upgrades as impacting electricity spending multiplicatively and audits as reducing uncertainty on that multiplier. In practice, audit programs convey information in a variety ways. Some programs may report a percentage savings for individual end uses affected by specific upgrades (e.g., efficient windows decrease heating costs by X%), while others may report expected savings in levels of dollar or energy units (e.g., efficient windows reduce heating costs of the average household by $Y). Some programs may simply allow households to learn about energy efficiency options by recommending a list or ranking of the upgrades likely to yield the greatest savings. In each of these cases, some percentage savings $(1 - \alpha)$ is implied in the information presented. Our simple formulation captures a wide variety of these cases. Even in cases where the audit report only provides a ranking of upgrade options, households still have less uncertainty over potential savings after the audit. We model this in a stark way by having the audit reduce uncertainty in α completely.

The model's combination of uncertainty in average usage and uncertainty in energy efficiency savings is particularly well suited to cases in which the household doesn't know how its own usage patterns or behavioral habits interact with the given durable goods stock. We can think of these behaviors or usage patterns as part of the household's "type" which are captured in μ. Heating, cooling, and lighting are all good examples for which the model applies particularly well. Appliances such as refrigerators whose energy savings have little variability and are primarily determined by product specifications (less uncertainty in μ), and for which the the savings from replacement are easily calculated (less uncertainty in α), are not as well described by the model. In that case, providing information in levels of savings will likely provide more information than percentages, as is done with energy star savings presented in dollars.

3.1.1 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 upgrade 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 upgrade. 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 upgrade are fixed over time. Previous studies present evidence that a sizable portion of the fixed costs 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 installing an energy efficiency upgrade 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. While time varying fixed costs can be applied to our model, the resulting insights surrounding the impacts of uncertainty would be unchanged.

Finally, many energy efficiency upgrades also improve comfort in the home, such as improved insulation that reduces drafts.[14] An explicit model of the comfort production function is beyond the scope of this paper, but our simple formulation captures many of the key trade-offs when energy efficiency upgrades increase net utility either by improving comfort or reducing energy expenditures. An analogy to our framework is that households are uncertain about their level of comfort when they move into a new home, and are uncertain about the comfort improvements that energy efficiency upgrades can provide. Living in the home for a longer period reduces the first source of uncertainty, and audits can reduce some of the second source of uncertainty by identifying upgrade options that are likely to achieve the largest improvements. Ultimately, as long as upgrades increase net utility and audits reduce uncertainty over the gains from an upgrade, then our model applies.

3.1.2 Model Timing

Figure 2 shows the timing of the model. Initially, we assume a household must get an audit before installing an upgrade. This reflects the structure of the audit program introduced

[14]Increased comfort would also arise if an energy efficiency improvement leads to a rebound effect in which the household's consumption of energy services (e.g., cooling or heating) increases.

in the preceding section. In Section 4, we explore the impacts of allowing households to install subsidized upgrades without first receiving an audit.

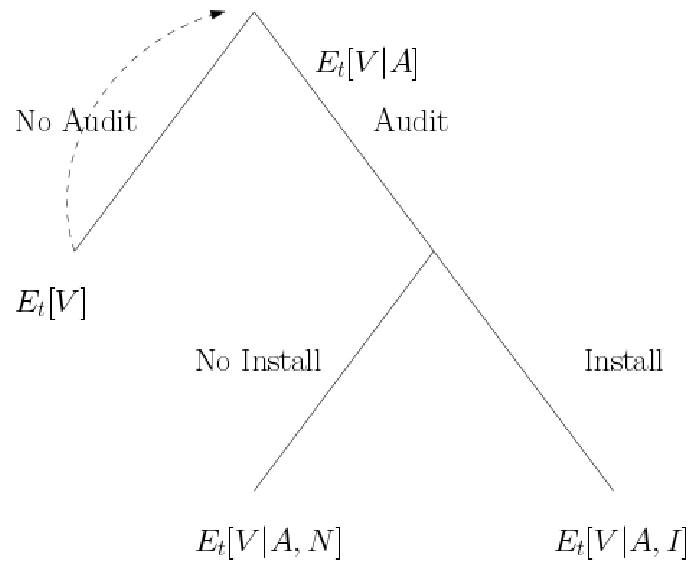

Figure 2: No audit, audit and install upgrades decision tree for each time period.

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 install an upgrade (I) or not install an upgrade (N). For simplicity we assume households make a "yes" or "no" upgrade decision in the same time step as when they get an audit. A household installs an upgrade 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 upgrade, is larger than the expected value of not doing so. Thus, expectations over savings from making an upgrade 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 upgrade decisions.[15]

[15]An alternative would be to fully model audits and upgrades as an exercise in option value. Given our assumption of unbiased priors over mean energy use and linear utility, the gains of modeling fully forward looking households are low.

3.2 Audit and Upgrade Decisions

We proceed by comparing expected value functions at each node in the decision tree shown in Figure 2. We start with the decision to install an upgrade post-audit. We then backward induct to the audit decision.

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

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

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

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

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 installing an upgrade and not, given their beliefs:

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

Thus, households install an upgrade 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 installing an upgrade for any given combination of upgrade cost (κ) and upgrade savings $(1-\alpha)$. Once an audit is performed and α is known, rearranging the expression in (4) provides the *ex post usage threshold*:

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

such that it is rational to make an upgrade if $m_t(\theta) \geq \bar{m}$.[16] In sum, these modeling assumptions have the intuitive implication that households are more likely to make upgrades 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 upgrade or not (realized α is below or above the critical value). Combining equations (2) and (3), the expected value function with an audit is:

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

Equation (5) shows that the expected value of getting an audit includes the value if no upgrade is made, weighted by the probability that α is above the critical level $(1 - G(\bar{\alpha}))$, plus the value with upgrade savings integrated over the support of α's in which an upgrade is made. For convenience we define the mean α conditional that an upgrade would be installed 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

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

This is the present value of income, less the fixed audit cost, less the expected fixed

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

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

upgrade cost, less the expected energy bills. These expected energy bills incorporate the probability of installing an upgrade and the expected savings if an upgrade is made. Note that if the probability of making an upgrade, $G(\bar{\alpha})$, 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 upgrade 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

$$E[V_{t+1}] = \underbrace{\frac{w}{1-\delta}}_{\text{present value of wealth}} - \underbrace{A \cdot \Big(1 - H(\widetilde{m})\Big)}_{\text{ex ante expected audit cost}} - \underbrace{\kappa \cdot G(\bar{\alpha})\Big(1 - H(\widetilde{m})\Big)}_{\text{ex ante expected upgrade cost}}$$
$$- \underbrace{\frac{p}{1-\delta}\Bigg(1 - G(\bar{\alpha}) \cdot (1 - \widehat{\alpha})\Bigg)(1 - H(\widetilde{m}))\widehat{\mu}_H}_{\text{expected bills if audit, given uncertainty over upgrade}} - \underbrace{\frac{p}{1-\delta} H(\widetilde{m})\widehat{\mu}_L}_{\text{expected bills if don't audit}} \quad . \tag{7}$$

This is the present value of wealth net of what may happen with audits, upgrades, 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 upgrade cost) accounts for the possibility of making an upgrade 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.

The audit decision therefore depends on the sign of the following expression:

$$E[V_t(audit)] - E[V_t(delay)] \gtreqless 0 \Longrightarrow$$
$$\frac{\delta}{1-\delta}\Big[\underbrace{\Big(1 - G(\bar{\alpha}) \cdot (1 - \widehat{\alpha})\Big)(1 - H(\widetilde{m}))p\widehat{\mu}_H + H(\widetilde{m})p\widehat{\mu}_L}_{\text{Expected future bills if delay, potentially audit later}} - \underbrace{\Big(1 - G(\bar{\alpha}) \cdot (1 - \widehat{\alpha})\Big)pm_t}_{\text{Expected future bills if audit now}} \Big]$$
$$+ \underbrace{G(\bar{\alpha})(1 - \widehat{\alpha}) \cdot pm_t}_{\text{Expected savings this period}} - \underbrace{\Big(A + \kappa \cdot G(\bar{\alpha})\Big) \cdot \Bigg(1 - \delta\Big(1 - H(\widetilde{m})\Big)\Bigg)}_{\text{Expected change in fixed costs}} \gtreqless 0. \tag{8}$$

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 upgrading 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 upgrade. 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 upgrade). In addition, the household accounts for the fixed costs from auditing now with certainty and installing an energy efficiency upgrade with probability $G(\bar{\alpha})$ versus potentially auditing and upgrading in the next period with probability $(1 - H(\widetilde{m}))$.

Finally, there are several theoretical implications from the model described above which can be tested empirically. In an Appendix, we describe these testable implications and present empirical evidence from the utility-run energy efficiency program described in Section II. However, we don't view these testable implications nor the empirical evidence consistent with them as the central contribution of the paper. Rather, we view the theoretical framework's implications for in home energy audit policy efficiency as the central contribution.

4 Policy Implications

In this section, we examine the theoretical model to assess welfare implications of programs subsidizing energy efficiency audits and upgrades. First we highlight that uncertainty surrounding the true payoffs from investing in energy efficiency can lead households to make upgrades even though they would not with full information. Put another way, a household might make an investment because it has positive net present value (NPV) in expectation, but from an ex post perspective (e.g. with full information) it was a mistake.

Second, we demonstrate that these ex post investment 'mistakes' can be exacerbated by typical energy efficiency audit and upgrade subsidy programs. This second point is a core contribution of the paper: while it is well known that a large share of subsidies go to inframarginal or 'non-additional' program participants, our analysis shows that even the households that are technically 'additional' to the program may not be economically efficient when assessed by a fully informed social planner. In a subsequent subsection, we calibrate the model to data from a medium sized MSA to assess what percentage of audit adopters might be making "full information mistakes".

4.1 Uncertainty, Investments and Household Types

Our model adds uncertainty around the benefits a household receives when they make an energy efficiency upgrade. A model without this type of uncertainty has two kinds of households: those that make upgrades, because the NPV is greater than zero, and those that do not, because the NPV is less than zero. Introducing uncertainty, we can divide households into four categories that depend not only on the sign of the NPV of an upgrade, but also on the sign of the household's expectation of the NPV of an upgrade. In particular, the four categories are: (1) households that make upgrades (i.e. positive expected NPV) that also have a positive NPV under full information, (2) households that correctly choose not to make upgrades (i.e. negative expected NPV) that would have negative NPV under full information, (3) households that delay making upgrades (negative expected NPV) that would have positive NPV under full information, and (4) households that make upgrades which are inefficient and don't pass a full information benefit-cost test (i.e. expected NPV positive, full information NPV negative). Highlighting the existence of these final two groups – types (3) and (4) above – is the core contribution of this paper. Delaying upgrades and making upgrades that don't pass the full information benefit-cost test is problematic because households often make upgrades shortly after moving in to a new home when their information about the home is lowest.

Which of the four groups a given household falls into depends crucially on their home's true mean energy usage and their prior beliefs about its energy usage. For example, if a household believes their home is an "energy hog" based upon observable characteristics but in fact it is not, they they might make an energy efficient investment that has a negative NPV. Conversely, if a household believes their home is very energy efficient based upon observable characteristic but it turns out to be an energy hog, they are likely to not immediately make an energy efficiency investment but would after learning about the true average electricity consumption of their home. We discuss below that the first case is more troubling than the second from a welfare perspective.

Figure 3 plots an example joint distribution of prior beliefs and true mean usage. At the time of move in, a household would be represented by a point somewhere on the X-Y plane. For example, a homeowner who believes the energy efficiency of their new purchased home is low would have a high prior for the mean of their electricity bill (e.g., be represented on the east side of the Figure). If the home turned out to be energy efficient and have low electricity usage, that household would be a point in the southeast part of the Figure. The overall distribution of all homes and homeowners at time of move in is represented by the arbitrary density along the Z-dimension of the Figure.

Recall, the theoretical model shows there is a threshold level of actual mean usage for which it is efficient for a household to make an energy installation. The solid black lines in Figure 3 represent this *full information energy usage threshold* ($\bar{m}$) which is the cutoff mean usage above which installing an upgrade 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 install an upgrade. However, because of uncertainty in the true mean, households will only install an upgrade if their prior belief (m_t) about their mean is above this threshold.

Figure 3 visually highlights where the four categories of households fall based upon a homeowner's prior over mean electricity use and the homes actual electricity use:

1. $\mu > \bar{m}$, $m_t > \bar{m}$, privately optimal to upgrade: the household invests in energy efficiency and, ex post, it is privately optimal to do so. This is represented in Figure 3 by the area under the compound distribution in the dark grey box in the upper right corner labeled "Correctly Install".

2. $\mu < \bar{m}$, $m_t < \bar{m}$, privately optimal not to upgrade: the household does not invest in energy efficiency, which is ex post the privately optimal decision. This is represented in Figure 3 by the un-shaded region in the lower left.

3. $\mu > \bar{m}$, $m_t < \bar{m}$, delayed upgrade: the household does not invest in energy efficiency when, if they had full information, it would be privately optimal to do so. This is

represented in Figure 3 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 as priors update to the 45 degree line.

4. $\mu < \bar{m},\ m_t > \bar{m}$, NPV negative investment: the household invests in energy efficiency upgrades that would not be made under full information. The households that fall into the light gray box in the lower right of Figure 3 invest in efficiency upgrades that are don't have a positive expected NPV when assessing with full information. They install an upgrade right away because their prior belief is above the threshold. As beliefs converge to their true mean usage over time these households eventually learn that their investment was not privately optimal.[17]

There is one important distinction to make about the economic costs of the two additional categories, delayed upgrades and NPV negative investment, introduced by allowing for uncertainty. Delayed upgrades are not ideal but might not be too detrimental from a welfare perspective because the long run benefits of the investment are still realized. Delaying an efficiency upgrade by, say, six months, leads to marginally higher expenditures, marginally more emissions from electricity consumption and inconvenience after a new homeowner is "settled in" to their new home. However, NPV negative investments have the potential to be much worse for welfare. NPV negative investments tie up capital irreversibly which could otherwise be used for more profitable endeavors.

[17] 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

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

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 upgrade subsidy, or show that as uncertainty in μ declines the share of households making mistakes also declines.

4.2 Visualizing Install Subsidies

Now consider how the decision to install an upgrade is affected by an upgrade subsidy. To simplify, assume the savings parameter (α) is known – i.e. we assume an audit has already occurred. Subsidizing installations reduces the needed expected monthly savings required for an installation to have positive expected NPV.

Figure 4 shows what happens visually when the install subsidy increases. By reducing the private cost of installing an upgrade, the subsidy reduces the upgrade threshold $\bar{m}$, which we depict in Figure 4 as a reduction in the upgrade threshold from $\bar{m}_0$ to $\bar{m}_1$. For simplicity, assume that the level of the subsidy is 'optimally' set such that it aligns the ex post privately optimal and ex post socially optimal choices. That is, if the net private benefits of an investment are positive with the subsidy, the net social benefits will also be positive (and if the net private benefits are negative with the subsidy, the net social benefits are negative). That is, it is privately and socially optimal for a household to perform an upgrade if and only if $\mu > \bar{m}_1$. With the lower upgrade threshold, a subset of the type (2) households, for whom it was privately optimal not to upgrade without the subsidy, now install an upgrade. Of these, only the households in the thickly cross-hatched area labeled "Target Additionality" become type (1) households for whom it is now privately and socially optimal to upgrade. 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.

As Figure 4 shows, however, a potentially significant share of the technically additional households may not be economically efficient. For example, the remaining households that did not upgrade 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 $\bar{m}$. In addition, some of the type (3) households are "pulled forward" into installing an upgrade earlier than they otherwise would have. These are the households in the medium gray, lightly hatched area above the "Target Additionality" region. It is privately and socially optimal for these households to install an upgrade, 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 types of new adopters are technically "additional" 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 upgrades 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 upgrades 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.

4.2.1 Visualizing Audit Subsidies

Our model also highlights that 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 upgrades without fully informing the investment decision. In appendix A.2, we show that the *full information 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 5, with $\widehat{\bar{m}}$ depicted with the dashed black line and the full information $\bar{m}$ depicted with the solid black lines. Audits induce more upgrades because they reduce the upgrade threshold through the resolution of uncertainty in α. Doing so may again "pull forward" a share of households that were delaying their upgrade installation due to priors about electricity usage being lower than actuals (the shaded region between the solid and dashed lines in Figure 4). 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 when they would not with full information.

4.3 How Important Is This?

To explore how prevalent ex post investment mistakes and delay mistakes are likely to be, we conduct a simulation exercise using our household-level data to parameterize the theoretical model. This is a descriptive exercise intended to give a rough sense of the frequency of the four household types depicted in Figure 3 and the three "additionality" categories depicted in Figure 4. To produce our simulation results, we first parameterize the compound distribution of a household's prior for their mean energy consumption conditional on the observed characteristics of the home ($m(\theta)$) and their true mean energy consumption (μ) – i.e. the distribution displayed in Figure 3.

For this simulation, we first need an estimate for priors about homes' energy consumption. These parameters are calculated from our three-year sample of household-level billing data. For example, 1,200 kWh/month is approximately the grand mean across individual households' mean monthly consumption, and 572 is the cross sectional sample standard deviation of household mean monthly consumption. Hence, we start with a normal distribution for $m(\theta)$ with a mean of 1,200 kWh/month and a standard deviation of 572 kWh/month.

For the conditional distribution of μ, we assume a normal distribution and that priors are unbiased and set $\mu = m(\theta) = 1200$.[18] The standard deviation of μ for a given $m(\theta)$ measures the degree of household uncertainty dictates the speed with which households update their priors. For example, if uncertainty around the true mean electricity consumption of a home is zero, a new homeowners updates after a single period.

To quantify the uncertainty in a household's expected energy consumption using the household-level data, we must strip away the variation in energy use that is driven by observable household characteristics (θ), leaving only the variation driven by unobserved home efficiency and household preferences. To do so, we run a cross-sectional regression of household-level mean monthly consumption over the three-year sample on a flexible func-

[18]To operationalize this we discretized the support of $m(\theta)$ into many small intervals and weighted the density of μ within a given $m(\theta)$ interval by the unconditonal density of $m(\theta)$ over that interval.

tion of the number of bedrooms, number of floors, the home age, and square footage.[19] Intuitively, the dependent variable in the regression – the sample mean of a household's energy consumption – is an estimate of the true mean μ. The fitted value of the regression is an estimate of the prior for mean monthly consumption $m(\theta)$ based only on observables. The variance of the residuals from the regression provides an estimate of the uncertainty that new occupants of the home would face about their mean usage. The standard deviation of these residuals is 542 kWh/month, which we use to parameterize the distribution of μ given $m(\theta)$.

We normalize this uncertainty by dividing by the grand mean of monthly energy use, resulting in a coefficent of variation (CofV) of 0.45. We consider this to be an upper bound on household uncertainty because the observed usage in the sample includes both intrinsic uncertainty and variation due to unobserved preferences for energy services. We therefore use the observed coefficient of variation based upon the regression approach, 0.45, as an upper bound and then perform sensitivity analysis using a coefficient of variation in increments between zero and 0.45.

In order to calculate the upgrade cutoff, $\bar{m}$ from Figure 3, we assume an electricity price of \$0.10/kWh, an eight percent annual discount rate, an α of 0.9 (i.e. ten percent energy savings) and an upgrade cost of \$4,000 without the \$500 subsidy.[20] With our estimate of the upgrade cutoffs, we can quantify the mass in each region of Figure 3 (i.e. "correctly install", "correctly don't install", "don't install but should", and "install but should not") for different levels of the coefficient of variation representing household uncertainty in mean energy use.

Figure 6 shows how uncertainty in mean energy use impacts the prevalence of ex post investment mistakes and delay mistakes. Consistent with the upgrade subsidy provided in our empirical setting, we also evaluate a \$500 subsidy for energy efficiency upgrades. Uncertainty clearly increases the prevalence of these ex post suboptimal decisions. Panel (a) shows that households making an ex post investment mistake, as a share of households that invest in energy efficiency upgrades, is large and increasing in household uncertainty.

[19] A detailed summary of the regression model is provided in Appendix C.

[20] Ten percent energy savings is in roughly what can be achieved by extensive duct sealing or window replacement. Although duct sealing is likely less than \$4,000, window replacement costs somewhat more than this.

Recall, absent any uncertainty, i.e. a CofV of zero, no households would make an ex post investment mistake (all investments would be privately optimal). However, even with low levels of uncertainty (e.g., CofV of 0.1), our simulation suggests that 20% of upgrades being performed would not be privately optimal. This share increases to 40 percent when uncertainty is at the upper bound of our estimate for recent movers, i.e. a coefficient of variation in expected mean use of about 0.45.

A $500 subsidy increases the number of households that make upgrades. Although Panel (a) shows that this reduces the share of upgrading households whose decision was not privately optimal ex post, Panel (b) shows that it also dramatically increases the total number of households making an ex post investment mistake.[21] This occurs because the subsidy increases the number of households making an upgrade, and a large share of these additional upgrades will not prove to be privately optimal ex post.

Panel (c) shows that the share of non-upgrading homes who are inefficiently delaying is fairly low. Ultimately, this is due to the fact that the vast majority of households in our simulation, and in practice, do not perform upgrades. Because of this feature, the share of all households who make a delay mistake depicted in Panel (d) looks very similar to the share of non-upgrading households making a delay mistake in Panel (c). Importantly, however, the share of non-upgrading homes who are inefficiently delaying upgrades again increases dramatically with uncertainty and with subsidies.

[21]This is a feature of the distribution, as a larger mass of marginal households are in the "Target Additionality" region in Figure 4 than in the "Ex Post Mistake" region.

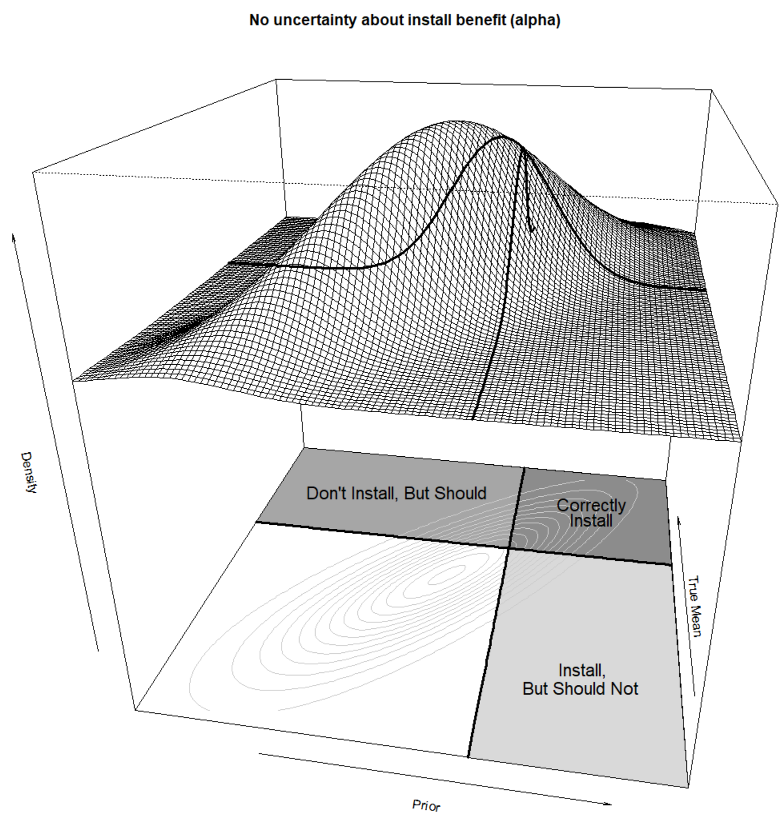

Figure 3: This figure delineates four categories of households. The share of households that install upgrades are those with a prior above the cutoff, given by the area under the compound distribution to the right of the vertical black line ("Install, But Should Not" and "Correctly Install"). The share of households that should install upgrades 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 an upgrade.

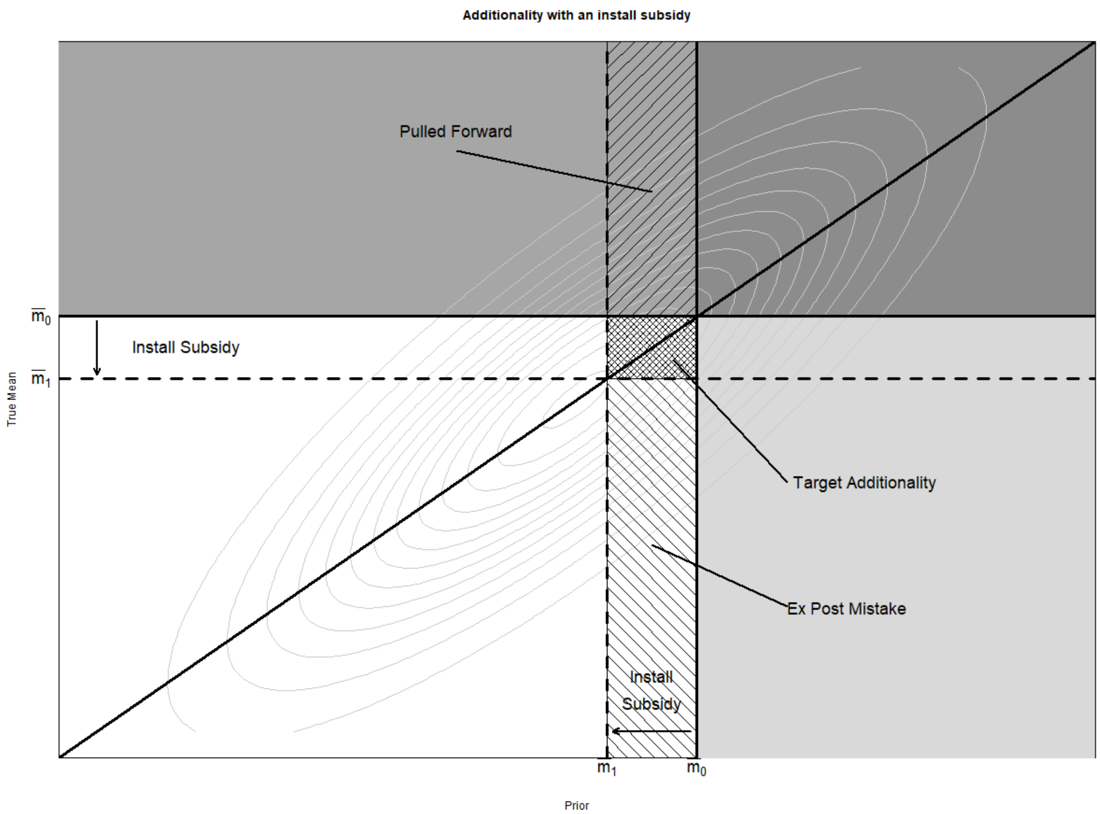

Figure 4: A subsidy reduces the critical value $\bar{m}$ for households to install an upgrade. The share of households that are "additional" to a subsidy policy are the area under the compound distribution between the vertical dashed black line ($\bar{m}_1$) and the vertical solid black line ($\bar{m}_0$). 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 installed an upgrade without the subsidy ("Pulled Forward"). The figure depicts the case when α is known.

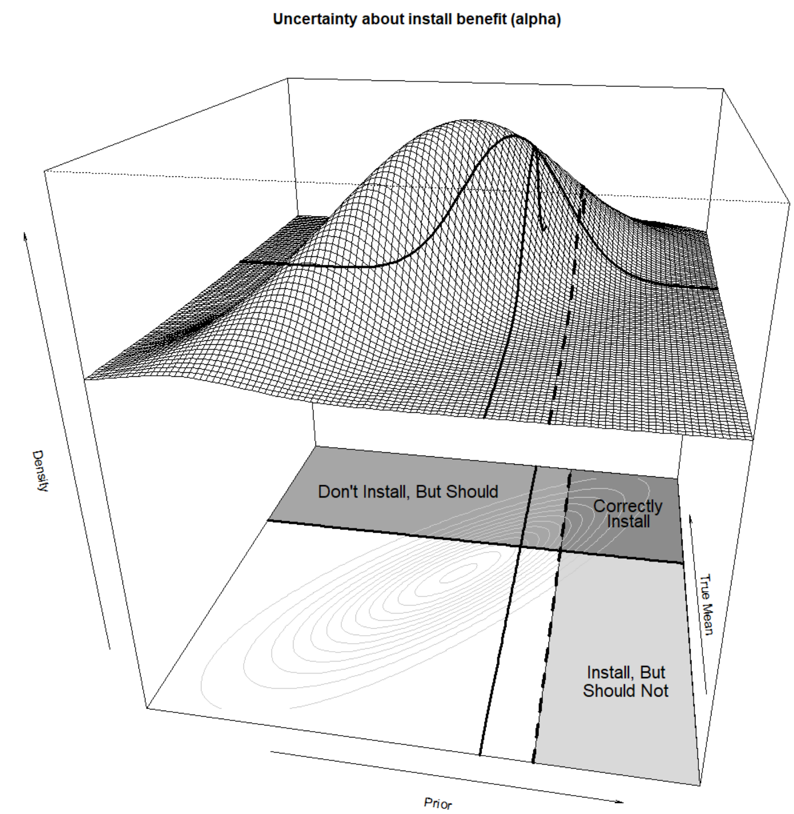

Figure 5: The share of households that install upgrades when α is uncertain is the area under the compound distribution to the right of the dashed black line ("Install, But Should Not" and "Correctly Install"). If uncertainty in α is removed, this cutoff moves to the left, to the solid black line. These additional participants are not necessarily economically efficient. The share of households that should install upgrades is the area under the compound distribution above the horizontal solid black line ("Don't Install, But Should" and "Correctly Install").

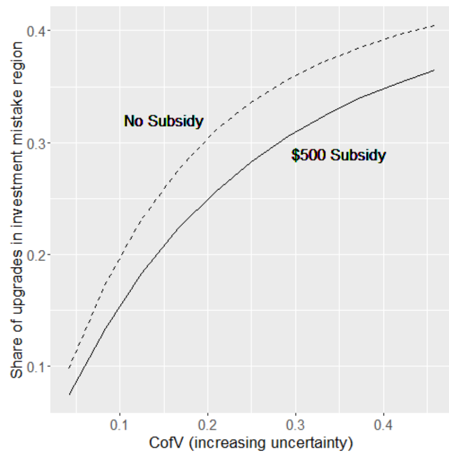

(a) Investment mistakes as a share of upgrades vs. uncertainty

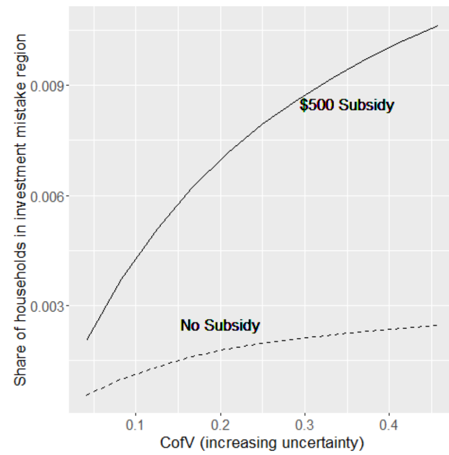

(b) Investment mistakes as a share of all households vs. uncertainty

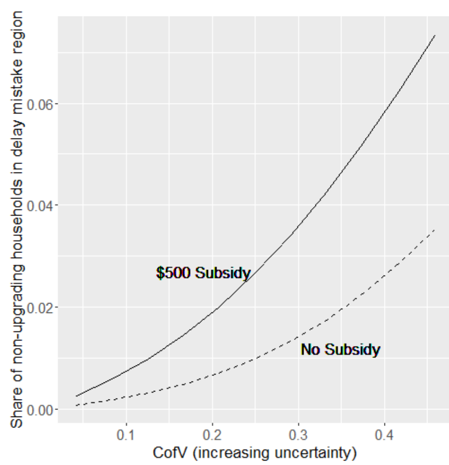

(c) Delay mistakes as a share of non-upgrading households vs. uncertainty

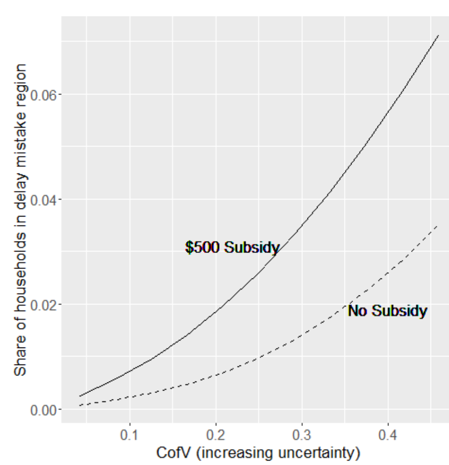

(d) Delay mistakes as a share of all households vs. uncertainty

Figure 6: This figure shows how the prevalance of investment mistakes and delay mistakes changes with household uncertainty and with an upgrade subsidy. Household uncertainty is measured as the coefficient of variation (CofV) in expected household electricity use. A $500 upgrade subsidy reduces $\bar{m}$, the upgrade threshold, and increases the share of households investing in energy efficiency upgrades.

We also use the simulation to decompose the additionality from the \$500 subsidy for energy efficiency upgrades. For simplicity, we again assume that the subsidy perfectly aligns the privately and socially optimal decisions – i.e. if it is privately optimal to perform an upgrade with the subsidy in place, it is also socially optimal to perform the subsidy.

Figure 7 presents a decomposition of the households that are "additional" to the \$500 subsidy, and shows how this decomposition depends on household uncertainty. We decompose the additional households into those that are unambiguously economically efficient adopters (Target Additionality), those that have their investments pulled forward in time, and those that are making investments which don't pass the full information benefit cost test. These correspond to the three additionality regions illustrated in Figure 4 between the vertical lines for $\bar{m}_1$ and $\bar{m}_0$. Among households that are additional to the subsidy program, Figure 7 shows that all of them fall into the target additionality category when there is no uncertainty. However, when the uncertainty is at our upper bound estimate (i.e. CofV of 0.45), the share of marginal participants that fall into the target additionality category falls to 24% while 40% of additional households are making investments that don't pass the full information benefit cost test and 36% are pulled forward.[22]

[22]Given the parameters in our simulation, about 80% of the households upgrading with the subsidy are additional and 20% upgrade even without the subsidy. These proportions do not change with uncertainty. Although the share of additionality is larger than estimates in previous literature, it is on the same order of magnitude.

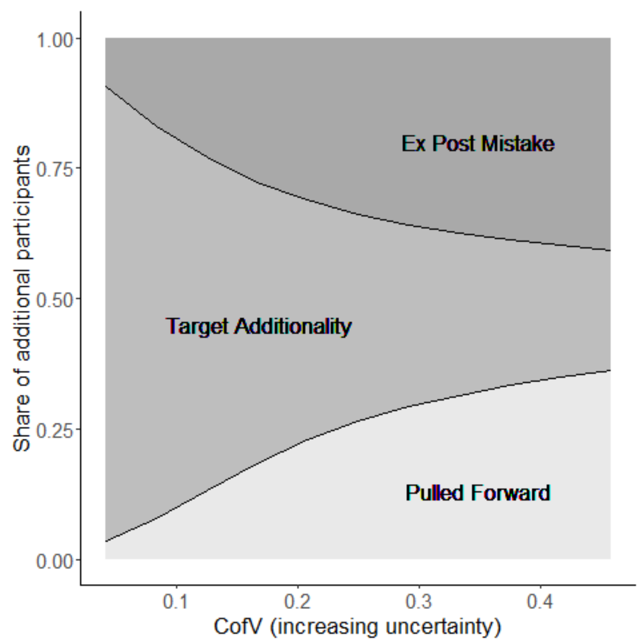

Figure 7: This figure decomposes the additional households – i.e. those that elect to perform energy efficiency upgrades only with the introduction of a \$500 upgrade subsidy – into those that are unambiguously economically efficient participants (Target Additionality) versus those that are making an ex post investment mistake or that are merely pulled forward. We show how the prevalence of each category changes with household uncertainty, which is measured as the coefficient of variation (CofV) in expected household electricity use.

4.4 Policy Implications

The insights provided by the theoretical model and intuition from the visualizations and simulation suggest that the social benefits provided by subsidizing energy efficiency upgrades may be smaller than previously thought due to uncertainty. 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 simply pulled forward in time. Moreover, we highlight that many of the additional participants may be making inefficient investments due to uncertainty – i.e. upgrades with costs larger than the stream of social benefits they will provide.

Importantly, policies often target financial incentives for energy efficiency upgrades directly towards households at the point of sale. For example, the federal government's energy efficiency mortgage program provides energy efficiency financing options to homeowners at the point of sale.[23] At first glance, targeting this mover margin with energy efficiency programs appears to have a great deal of scope. Not only do we observe that households are more inclined to participate immediately after moving in, 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 period.[24] However, our results suggest the need for a great deal of caution for targeting recent movers. If households have the most uncertainty surrounding the benefits from making energy efficiency upgrades immediately after moving in to a home, then targeting energy efficiency subsidies 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.

More generally, our analysis highlights that efficiency gains could be achieved by eliminating the uncertainty households face when deciding to make energy efficiency upgrades. This point is certainly something policymakers are aware of. In many States and mu-

[23] 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`.

[24] 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.

nicipalities, it is now a requirement that past energy bills be disclosed to home buyers, similar to how EPA mandates MPG and gas expenditures on sales of new vehicles. In some locations, it is now even a requirement that homeowners perform an energy audit and provide the findings to prospective buyers.[25]

Our results provide additional motivation surrounding pro-information point of sale policies. As we point out, a typical audit can provide information surrounding the energy savings that could be achieved by performing upgrades, conditional on some baseline usage pattern. Importantly, as Figure 4 highlights, it is not sufficient to simply remove the uncertainty surrounding the share of energy use that could be avoided (α) – or equivalently, the amount of energy that could be saved conditional on some baseline behavior and energy usage. It is also important to ensure that households have a precise understanding of their own baseline energy usage (μ) in their new home.[26] 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 empirical observation that households are most likely to participate in energy efficiency programs at the time they move into a home, this suggests 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.

5 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

[25]Recent work (Cassidy (2019), Myers, Puller and West (2019)) examines the impact of providing energy audit information at the point of sale in the Energy Conservation Audit and Disclosure (ECAD) ordinance in Austin, TX.

[26]Bill disclosures ultimately will not provide new owners a clear signal of their future baseline energy usage in a home. This is due to the fact that the past usage is dependent on the unobserved consumption patterns and behavior of the previous occupants.

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 investments that are NPV negative with full information – 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.

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.

Cassidy, Alecia. 2019. "How Does Mandatory Energy Efficiency Disclosure Affect Housing Prices?" *Available at SSRN 3047417*.

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.

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.

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.

Globus-Harris, Isla. 2019. "Estimating Free-Ridership in Energy Efficency Subsidy Programs." Working Paper.

Hassett, Kevin A, and Gilbert E Metcalf. 1992. "Energy tax credits and residential conservation investment." National Bureau of Economic Research.

Hassett, Kevin A, and Gilbert E Metcalf. 1995. "Energy tax credits and residential conservation investment: Evidence from panel data." *Journal of Public Economics*, 57(2): 201–217.

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.

Houde, Sébastien, and Joseph E Aldy. 2017. "Consumers' Response to State Energy Efficient Appliance Rebate Programs." *American Economic Journal: Economic Policy*, 9(4): 227–55.

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.

LaRiviere, Jacob, David Kling, James N Sanchirico, Charles Sims, and Michael Springborn. 2018. "The treatment of uncertainty and learning in the economics of natural resource and environmental management." *Review of Environmental Economics and Policy*, 12(1): 92–112.

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.

Myers, Erica, Steven L Puller, and Jeremy D West. 2019. "Effects of Mandatory Energy Efficiency Disclosure in Housing Markets." National Bureau of Economic Research.

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.

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:

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

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):

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

A.2 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 $E(\alpha)$:

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

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:

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

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 Testable Implications

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 B.1). First, recent movers have wide priors over their mean energy bills. As beliefs become more 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.

There is an alternative explanation for why recent movers may audit at higher rates not involving Bayesian updating: the effective cost (κ) of making an install could increase as homeowners spend more time in their homes. This explanation, however, offers no insight into 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. Our model can make additional testable 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 time in home increases. These predictions are driven by differential impacts of changes in uncertainty and the value of information as a function of observables in our data that would not arise from changes in the fixed costs of investment.

We now characterize how 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. The marginal influence of this information on audits and installations should decrease the more observations a household has about their mean energy use. We first show how information affects the priors about mean energy use over time, and then describe the impact of information on the value of delaying an audit or an installation one period.

In our model, each of these types of information 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 B.2):

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.[27] 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 B.3, we show that the partial derivative of (8) with respect to θ is given by

$$\frac{\partial E[V_t(audit)] - E[V_t(delay)]}{\partial \theta} = \left[\frac{p}{1-\delta} G(\bar{\alpha})(1-\widehat{\alpha})\Big(1-\delta+\delta H(\widetilde{m})\Big)+\delta h(\widetilde{m})(\tau+tr)A\right] m'_{t,\theta}$$
$$+\frac{\delta p}{1-\delta}\left[h(\widetilde{m})(\tau+tr)G(\bar{\alpha})\Big((1-\bar{\alpha})m_t-(1-\widehat{\alpha})\widetilde{m}\Big)+g(\bar{\alpha})\frac{\partial \bar{\alpha}}{\partial m_t}(1-\bar{\alpha})(1-H(\widetilde{m}))(m_t-\widehat{\mu}_H)\right] m'_{t,\theta}. \tag{A3}$$

Equation (A3) 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 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.[28] 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

[27] 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.

[28] As shown in Lemma 2, both terms in equation (A3) are multiplied by $m'_{t,\theta}$, which decline over time.

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*

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 B.4). 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.

B.1 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:

$$N_t = \int_{-\infty}^{\widetilde{m}} \Big(1 - H_t(\widetilde{m}; \tau + tr, \theta)\Big) \prod_{s=1}^{t-1} H_s(\widetilde{m}; \tau + sr, \theta) dM(m(\theta)) \tag{A4}$$

Equation (A4) 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, $(1 - H_t) > (1 - H_{t+1})$.

B.2 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.

B.3 Derivation of equation (A3)

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 (A3).

Simplifying (8) slightly:

$$\begin{aligned} E[V_t(audit)] - E[V_t(delay)] &= G(\bar{\alpha})(1-\widehat{\alpha}) \cdot pm_t \\ &\frac{\delta}{1-\delta} G(\bar{\alpha})(1-\widehat{\alpha}) H(\widetilde{m}) p \widehat{\mu}_L - \Big(A + \kappa \cdot G(\bar{\alpha})\Big) \cdot \bigg(1 - \delta\Big(1 - H(\widetilde{m})\Big)\bigg) \end{aligned} \tag{A5}$$

Differentiating (A5) with respect to θ gives

$$\begin{aligned} \frac{\partial E[V_t(audit)] - E[V_t(delay)]}{\partial \theta} &= G(\bar{\alpha})(1-\widehat{\alpha}) p m'_{t,\theta} + pm_t \frac{\partial G(\bar{\alpha})(1-\widehat{\alpha})}{\partial m_t} m'_{t,\theta} \\ &+ \frac{\delta p}{1-\delta} \left(G(\bar{\alpha})(1-\widehat{\alpha}) \frac{\partial H(\widetilde{m})\widehat{\mu}_L}{\partial m_t} + \frac{\partial G(\bar{\alpha})(1-\widehat{\alpha})}{\partial m_t} H(\widetilde{m})\widehat{\mu}_L \right) m'_{t,\theta} \\ &+ \delta\Big(A + \kappa \cdot G(\bar{\alpha})\Big) \frac{\partial(1-H(\widetilde{m}))}{\partial m_t} m'_{t,\theta} - \kappa\Big(1 - \delta(1-H(\widetilde{m}))\Big) \frac{\partial G(\bar{\alpha})}{\partial m_t} m'_{t,\theta} \end{aligned} \tag{A6}$$

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

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

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 (A6) yields

$$\begin{aligned}
\frac{\partial E[V_t(audit)] - E[V_t(delay)]}{\partial\theta} &= G(\bar{\alpha})(1-\widehat{\alpha})pm'_{t,\theta} + pm_t(1-\bar{\alpha})g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}m'_{t,\theta} \\
+\frac{\delta p}{1-\delta}\Bigg(G(\bar{\alpha})(1-\widehat{\alpha})\Big(-\widetilde{m}h(\widetilde{m})(\tau+tr)+H(\widetilde{m})\Big) &+ (1-\bar{\alpha})g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}H(\widetilde{m})\widehat{\mu}_L\Bigg)m'_{t,\theta} \\
+\delta\Big(A+\kappa\cdot G(\bar{\alpha})\Big)h(\widetilde{m})(\tau+tr)m'_{t,\theta} &- \kappa\Big(1-\delta(1-H(\widetilde{m}))\Big)g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}m'_{t,\theta} \qquad \text{(A7)}
\end{aligned}$$

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

$$\frac{\partial E[V_t(audit)] - E[V_t(delay)]}{\partial\theta} = \left[\frac{p}{1-\delta}G(\bar{\alpha})(1-\widehat{\alpha})\Big(1-\delta+\delta H(\widetilde{m})\Big)+\delta h(\widetilde{m})(\tau+tr)A\right]m'_{t,\theta}$$
$$+\frac{\delta p}{1-\delta}\left[h(\widetilde{m})(\tau+tr)G(\bar{\alpha})\Big((1-\bar{\alpha})m_t-(1-\widehat{\alpha})\widetilde{m}\Big)+g(\bar{\alpha})\frac{\partial\bar{\alpha}}{\partial m_t}(1-\bar{\alpha})(1-H(\widetilde{m}))(m_t-\widehat{\mu}_H)\right]m'_{t,\theta} \tag{A8}$$

B.4 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:

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

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.

B.5 Empirical Evidence for Theoretical Predictions

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.[29] 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.[30] 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

[29]The data was shared under a privacy agreement directly by the auditing agency.
[30]This data was publicly available from the county assessor in our study footprint.

end dates.[31]

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.[32] 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 sample period.[33] 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 C highlighting that energy use varies dramatically across premises, even after conditioning on observable characteristics. Observing large variation in household electricity use conditioning on a subset of observables is consistent with learning being important upon moving into a home. Its possible someone could move into a home that appears to be an "energy hog" based upon observables and make an install based upon expected savings only to discover the home was actually energy efficient and

[31]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.

[32]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 err on the conservative side and drop premises which do not have a clear address match across the two samples.

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

the install was a mistake. This is a narrative we explore in greater depth following the empirical tests of the theoretical model.

B.6 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 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. Recall, the top panel of Figure 1 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.[34]

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 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.

[34] 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.

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 A1 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.

B.7 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 an inspection and upgrades performed on the home.[35]

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 **??**). Instead, the 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 **??**'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

[35]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.

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:[36]

$$\text{Audit}_{i,t} = \alpha_i + \gamma_m + \beta_1 \cdot \text{Elec}_{i,t} + \beta_2 \cdot \text{Gas}_{i,t} + \theta_1 \cdot \text{Elec}_{i,t-1} + \theta_2 \cdot \text{Gas}_{i,t-1} + \varepsilon_{i,t}. \quad \text{(A9)}$$

In the model specified by Eq. A9 above, $\text{Audit}_{i,t}$ is an indicator variable which equals one during the month t in which household i receives an energy audit.[37] $\text{Elec}_{i,t}$ 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.[38] $\text{Elec}_{i,t-1}$ and $\text{Gas}_{i,t-1}$ are the average daily 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.[39]

The key coefficients of interest from Eq. A9 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 A2 presents estimates from the general model specified by Eq. A9. The first

[36]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.

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

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

[39]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.

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. A9 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. 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. A9, 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 A2 present the estimates of Eq. A9 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.

B.8 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.[40] We classify a household as a mover if an audit was received anytime after the move-in date during our three year sample. Figure A4 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.

While it is possible factors not directly captured by our theoretical model could explain the fact that audits were more likely among movers compared to non-movers, the differential impact by home age is more challenging. For example, the fixed cost of performing an audit may increase with time spent in a home. In this setting, we would expect a higher likelihood for audits among the group of movers as opposed to non-movers.[41] Importantly, while differences in the fixed audit costs would contribute to the level difference between the likelihoods displayed in Figure A4, it would not explain the fact that the likelihood of an audit varies with home age early in a homeowner's tenure but not once they have resided in the premise for multiple years. We take this empirical pattern as strong evidence that, during the period immediately following a move, households use observable premise characteristics (e.g., home age) to inform their beliefs about the potential returns to audits and energy efficiency upgrades. Indeed, insofar as costs of making energy efficiency

[40] 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.

[41] This pattern would also be driven by harvesting. If the households that were more likely to seek out an audit have already had one, then the set of non-movers should be more heavily comprised of households that are less likely to choose to receive an audit.

installations are lower at time of move in, it increases the value of providing information on a home's actual average energy usage to that home's purchaser at time of move in for making installations that pass a benefit-cost test.

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 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:

$$\text{Audit}_i = \alpha + \beta \cdot \text{Year}_i + \theta \cdot \text{Year}_i \cdot \text{Mover}_i + \boldsymbol{\phi} \cdot \boldsymbol{X}_i + \varepsilon_i. \tag{A10}$$

In the model specified by Eq. A10, Audit_i is an indicator variable which equals one for each premise i that is audited at any point during our three year sample.[42] 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. A10, 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 A3 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 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

[42] Note, in contrast to the model specified by Eq. A9, which explored the timing of audits, the audit indicator now only varies across premises, not across time.

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 A5 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. A10 excluding the homes that are audited within the first two months of moving (results reported in column four of Table A3). Consistent with Figure A5, 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).

B.9 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. Moreover, observing a non-trivial share of audited households electing not to perform any subsequent upgrades (1/3 in our data) provides strong evidence that the information households receive through the audit is in fact being used. If households instead believed that the audit provided no reliable information, then only households that were determined to make efficiency upgrades would pay for the audit.

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 2 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 our 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 A6 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 A7 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 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:

$$\text{Install}_i = \alpha + \beta \cdot \text{Year}_i + \boldsymbol{\phi} \cdot \boldsymbol{X}_i + \varepsilon_i. \tag{A11}$$

In the model specified by Eq. A11, 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. A11 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 A4 present the point estimates of β with and without the observable controls. Consistent with the pattern displayed in Figure A6, 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 A4 include point estimates of the difference between β from Eq. A11 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.

C Measuring the Magnitude of Electricity Variability

To size how much uncertainty of electricity usage at time of move in matters, we regressed each household's mean monthly electricity use on a set of observables such as number of bedrooms, number of floors, a quadratic function of home age, and a cubic function of square footage. Specifically, we estimate:

$$\overline{usage_i} = \alpha + \Sigma_{s=1}^{2} yearbuilt_i^s \beta_s + \Sigma_{s=1}^{3} squarefeet_i^s \delta_s + \#bedrooms_i \psi + \#floors_i \rho + \epsilon_i \tag{A12}$$

All coefficients are statistically significant at well above the 1% level in this regression. We then take the predicted value of $\widehat{\overline{usage_i}}$. We then set this predicted mean usage to priors as a function of observables from the theoretical model: $\widehat{\overline{usage_i}} = m(\theta)$. Accordingly, the dependent variable in the regression — the mean usage observed by the occupants — is an estimate of the true mean μ. The variance of the residual from the regression provides an estimate of the uncertainty that new occupants of the home would face about their mean usage.

The average monthly usage in our estimating sample is roughly 1208.5 KWhs/month and the standard deviation around that from the regression (the square root of the estimated variance of the regression residual) is roughly 542. Therefore, the residual variance divided by the grand mean of household electricity use in the sample, results in a coefficient of variation of about 0.45. We take this to be the upper bound of calibrated priors of a home's electricity consumption relative to the home's realized mean electricity consumption.

Measuring variation around predicted usage will overstate the coefficient of variation so we assess how different levels of prior precision over a home's electricity consumption will impact ex post inefficient investments and pulled forward investments. In an ideal world we would observe the electricity consumption for every home in the dataset for a homeowner with "modal" preferences for climate comfort, lighting, price sensitivity, etc. Of course we don't observe that ideal. In our data, observed variation in a home's electricity consumption is a combination of both observed and unobserved fixed home charateristics and homeowner preferences. As a result we take the observed coefficient of variation based upon the regression approach, 0.45, use it as an upper bound and then perform sensitivity analysis between a coefficient of variation between zero and 0.45.

C.1 Another View of 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.

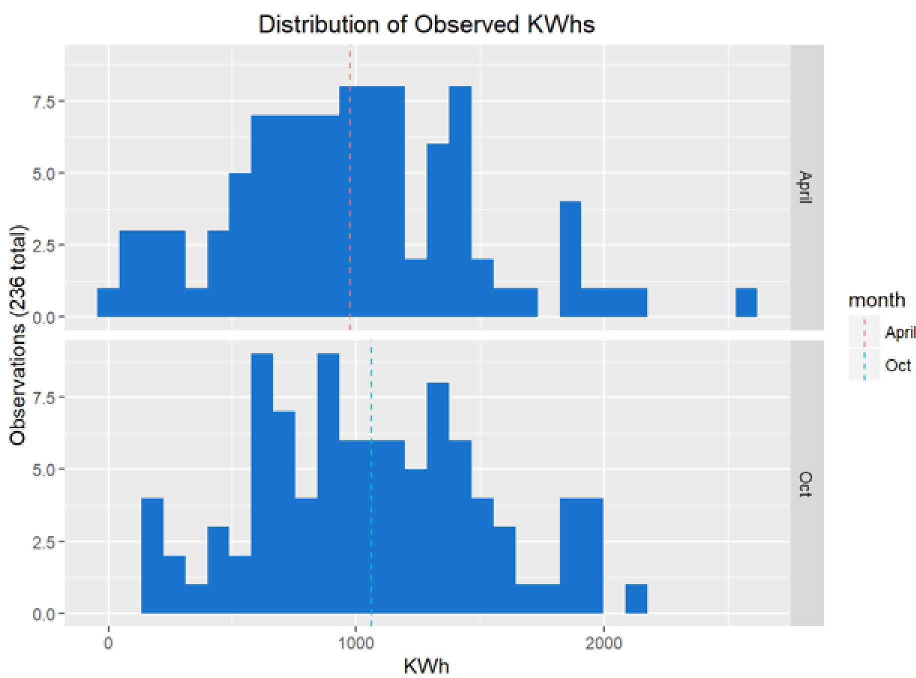

Figure A1: Notes: Observed April and October 2012 electricity usage.

Figures A1 and A2 show observed and predicted electricity usage of the 236 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

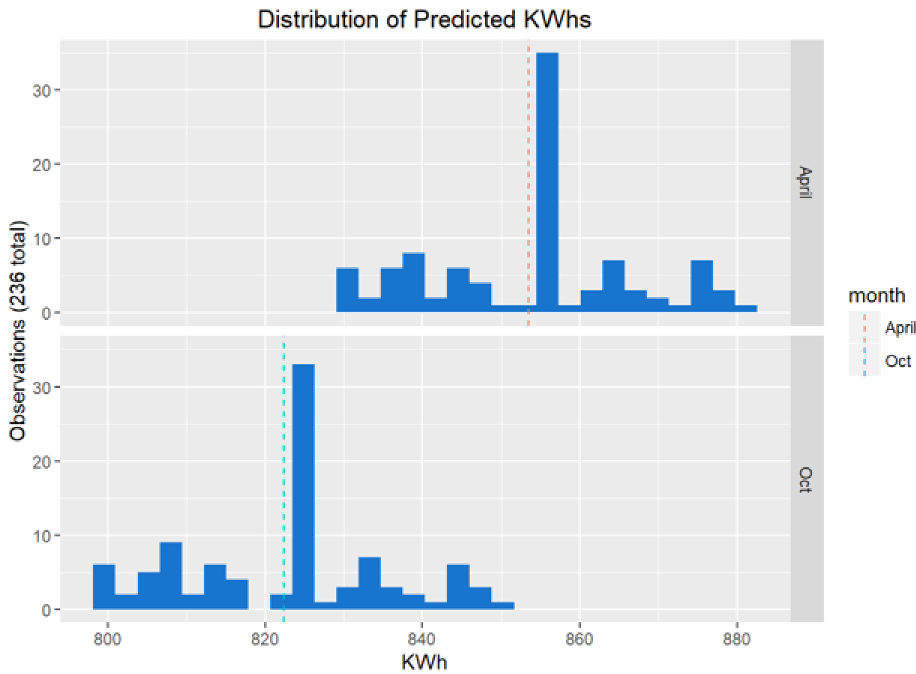

Figure A2: Notes: Predicted April and October 2012 electricity usage.

observables.

D Appendix Tables and Figures

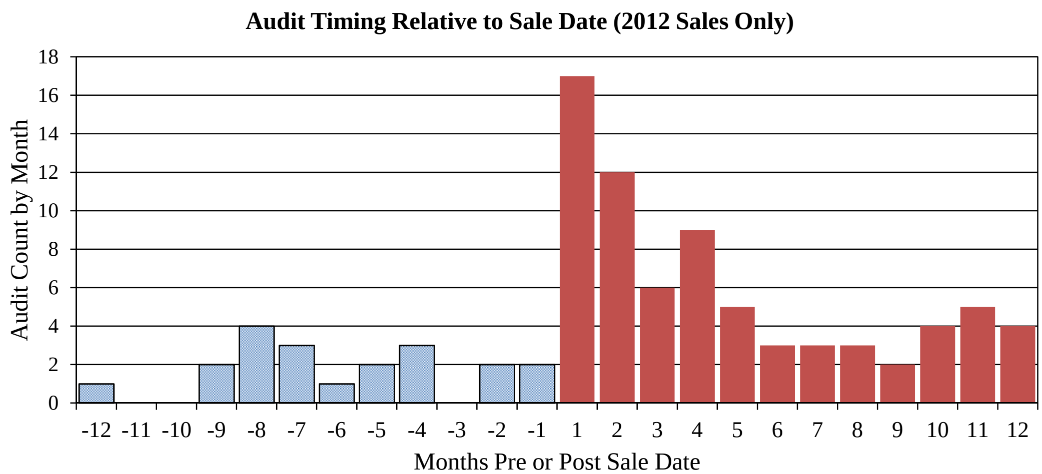

Figure A3: Audit timing relative to home sale for all homes sold in 2012 which also had an audit.

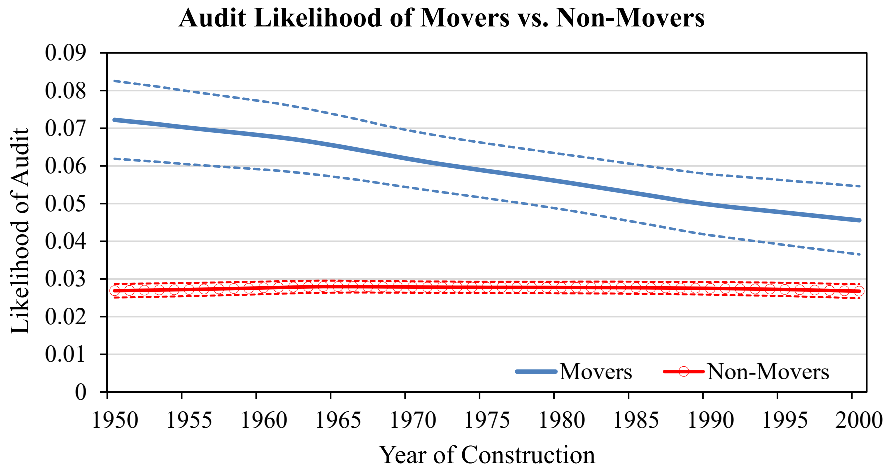

Figure A4: Likelihood of audit by home construction year conditional on a home being sold (Movers) versus not sold (Non-Movers) in our sample.

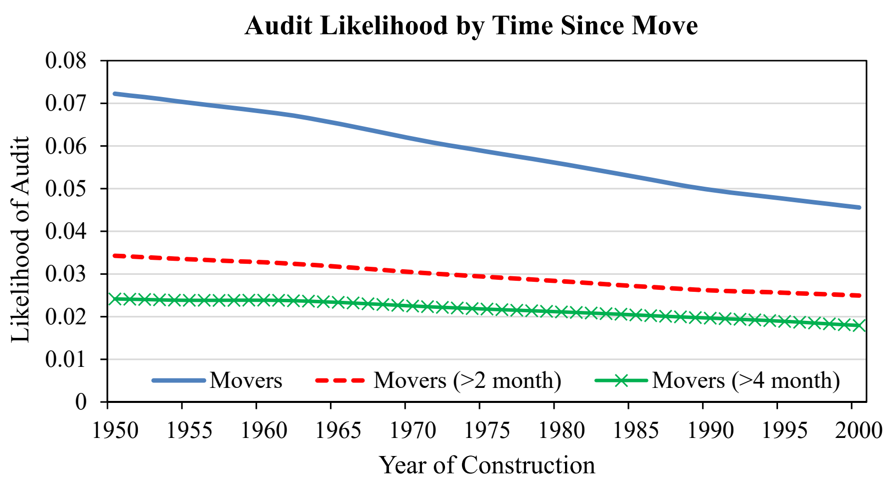

Figure A5: Likelihood of audit by home construction year conditional on a home being sold (Movers) as a function of when the audit occurs relative to the home sale. The Figure shows that observable characteristics like home age decrease in importance in predicting audit probably as time after sale increases for sold homes that have an audit.

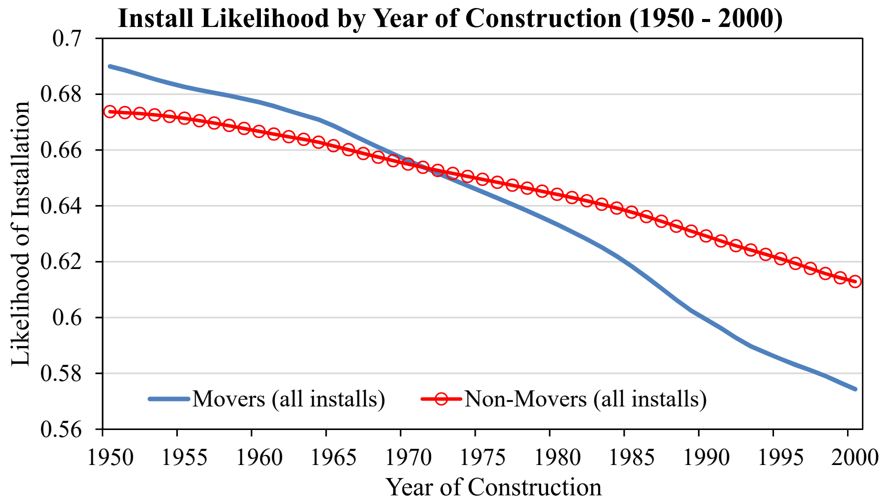

Figure A6: Install likelihood conditional on having an audit for home that are sold (Movers) versus not sold (Non-Movers) by year of construction. The Figure shows movers are much less likely to make installations in newer homes conditional on having an audit and more likely to make installs in older homes.

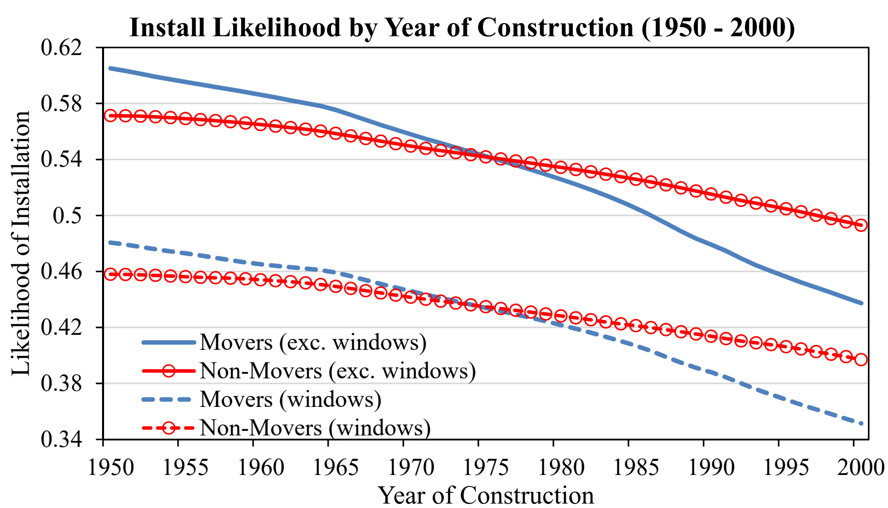

Figure A7: Notes

Table A1: 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 A2: 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 A3: 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 A4: 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.