Durable Goods and Long-Run Electricity Demand: Evidence from Air Conditioner Purchase Behavior

David Rapson*
Department of Economics, University of California, Davis

October, 2011

Abstract

I estimate a dynamic structural model of demand for energy-intensive durable goods, focusing on air conditioners. Expectations and uncertainty are explicitly incorporated into the model of consumer behavior. This is crucial to accurately estimating the timing of durable good purchases and derived long-run energy demand. I calculate important, policy-relevant demand elasticities, rare in the literature. Results indicate that there are interactions between the intensive and extensive margins, and that consumers are forward-looking and value energy efficiency. I reject the hypotheses of myopic and (more weakly) naïve expectations models of consumer behavior in favor of rational expectations.

*I thank Marc Rysman, Scott Carrell, Randall Ellis, Matthew Harding, Jordi Jaumandreu, Chris Knittel, Laurence Kotlikoff, Michael Manove, Robert Margo, J. Bradford Rice, Jeremy Smith, and Robert Stavins for their comments. Excellent research assistance during earlier revisions was provided by Tom Blake. This work has benefited also from input by seminar participants at Brown, the Federal Reserve Board, Harvard, HEC Montreal, IIOC, McMaster, Queen's, Stanford, UC Davis, and the UC Energy Institute. Discussions with participants in the Harvard Environmental Economics Program are also appreciated. All errors are my own.

1 Introduction

I estimate a dynamic discrete choice model of the demand for air conditioners, the most energy-intensive home durable good in the United States.[1] The goal is to provide a model for studying appliance choice and energy usage, and to improve our understanding of how long-run electricity demand is affected by consumer decisions on the intensive and extensive margins. Over time, air conditioners have exhibited large energy efficiency gains and even larger real price decreases.[2] In such a market, a fully-informed and rational consumer will form expectations about future changes in these key product attributes, and time her purchase to maximize the present value of the associated stream of net benefits. I explicitly test this rational expectations hypothesis against two alternate candidates of consumer behavior in dynamic settings: myopia, where consumers fully discount future utility, and naïve expectations, in which future realizations of key attributes are expected to remain unchanged relative to their current value.

Results from the model allow me to calculate the first estimates of the elasticities of demand for energy-consuming durables derived from a fully rational framework of behavior. Specifically, the estimated distribution of model primitives provides the necessary inputs to simulate the distribution of demand elasticities with respect to the durable's price, energy efficiency, and primary input cost (in this case, the price of electricity). These estimates help to alleviate a scarcity of empirical data points on the topic; however, their main value is to deepen our understanding of how consumers behave in this setting, and to use these insights to inform policy.

As is the case with most energy-consuming durables, the full social cost of use is not internalized by the user. External costs of energy use include the harm generated from carbon emissions (climate change) and criteria pollutants (health effects of poor ambient air quality). A common recommendation among economists is to impose a Pigouvian tax that aligns private and public costs. However, for reasons unrelated to economic efficiency, regulators are slow to embrace this solution. Instead, energy efficiency standards have been the most common tool that policy-makers use to reduce the external cost of energy use.[3] It is thus reasonable to seek to quantify the extent to

[1] In 2001, 16 percent of residential electricity in the U.S. was used for air conditioning (*Energy Information Administration*).

[2] Room air conditioners became 23 percent more efficient and 49 percent less expensive from 1987 to 2005, while central units became 27 percent more efficient and, depending on their size, between 14 to 49 percent less expensive.

[3] It is by no means the only approach. In some markets, like sulphur dioxide, cap-and-trade has been implemented and proven effective; in others, like those regulated by the Clean Air Act and its amendments, a more direct approach has been taken.

which these achieve their intended effect. A first step to understanding whether efficiency standards reduce energy usage (a necessary but not sufficient condition for economic efficiency), one must first understand how consumer demand responds to changes in energy efficiency, the relevant product attribute.

Existing studies examine the durable good demand response to electricity price fluctuations and energy efficiency standards using frameworks that only measure near-term outcomes. Dubin and MacFadden (1984) and Bento et al. (2009), for example, implicitly assume that the market for durables operates as a perfect rental market, allowing consumers to costlessly change attributes of their durables each period. In his seminal paper on consumer discount rates (also in the context of air conditioners), Hausman (1979) takes a more direct approach to estimating the tradeoff between short-run air conditioner operating costs and the cost of buying a more energy-efficient unit. His analysis is conditional on the decision to purchase. When we are interested in long-run outcomes, purchase timing (that is, the decision of whether to buy today or wait) may have important implications for the evolution of the stock of durables. Omitting these factors has the benefit of simplifying the analysis, but doing so risks creating an inaccurate representation of long-run effects of policies being considered. On the other hand, the added complexity comes at a cost: computation takes substantially longer than for reduced form approaches, and some relevant data must be discarded for computational parsimony.

An accurate long-run outlook of energy demand should use a dynamic framework of durable good purchases that encapsulates intensity-of-use (and the future stream of production costs that it implies) that varies with energy efficiency and energy prices. The model used herein builds on a framework developed in industrial organization and labor economics, drawing primarily on Rust (1987). Rust estimates the optimal replacement timing of bus engines, a decision that has several similarities to the timing of consumer durable good purchases. I extract dynamic content from appliance age in a repeated cross section of survey data, and extend Rust's model to incorporate large amounts of observed consumer heterogeneity. Individual level data on demand makes my empirical approach similar to Goldberg (1995), extended to include dynamic considerations.

This study has five main qualitative findings. First, I reject the hypotheses of myopic and, more weakly, naïve expectations models of consumer behavior in favor of rational expectations. Second, consumers exhibit higher (in absolute value) elasticities of demand with respect to energy efficiency than with respect to either the durable price or the price of electricity. While economic theory dictates that consumers should respond symmetrically to changes in energy efficiency and electricity prices, I allow these effects to differ in my model. This is consistent with other energy demand studies[4] and is motivated by the observation that consumers actually appear to respond differently to economics incentives on these margins. Third, increases in the energy efficiency of air conditioners raises air conditioner demand and lowers overall energy consumption. Fourth, increased electricity prices lower current energy consumption but have a statistically-weak negative effect on demand for air conditioner units. Fifth, lower air conditioner prices increase unit demand and slightly lower energy consumption because consumers replace old inefficient units with newer more efficient ones. Efficiency gains outweigh first-time purchasers and platform upgrades (from room to more energy-intensive central AC). These results are robust to a wide range of values of key parameter inputs.

The rejection of consumer myopia and high efficiency elasticities of demand combine to suggest that consumers are forward-looking. This deviates from the conventional view of consumers as myopic. If consumers did in fact discount the future entirely (or at a very high rate), they would be highly sensitive to the up-front durable price and less so to unit efficiency or electricity prices. This is because the latter two are a stream of future household production costs. However, results indicate that demand for air conditioners is own-price inelastic. Appliance price weakly influences the purchase decision, but not intensity-of-use thereafter. Electricity prices, on the other hand, strongly affect the intensity-of-use of durables, but less so the timing of purchases. While theory may lead to ambiguous predictions about the sign of the electricity-price elasticity of air conditioning unit demand, results here suggest that it is negative.

2 Related Literature

This paper contributes to the literature on the demand for energy-consuming durable goods, and determinants of change in residential energy use behavior. The methods used here draw upon an extensive literature on modeling dynamic discrete choices. In this section I will discuss how this paper relates to each of these areas of the literature.

The energy crisis of the 1970s focused economists' attention on energy demand. Early studies examined consumer discount rates, estimated short-run demand elasticities, and made initial attempts to formalize the relationship between appliance choice and usage. From the mid-1980s to the late

[4] E.g. Hausman (1979) and Bento et al. (2009)

1990s, interest in energy-related issues waned, but empirical methods and computing power have advanced. Recently, with energy having emerged once again as an issue of pressing concern, a new wave of economic research has been conducted. These newer studies have refined the approaches used earlier and in some case shifted focus to longer-run questions, including whether increases in energy prices induce technology improvements in energy efficiency via research and development, and the extent to which evolution of the stock of durables influences energy usage.

Of the earlier papers, Hausman (1979) and Dubin and MacFadden (1984) endure as seminal works. Hausman estimates the implied consumer discount rate associated with investment in energy-conserving appliances. The approach in this paper shares several important similarities to Hausman's, but they differ most starkly in the treatment of consumer expectations. Households in his model purchase air conditioners at an exogenously-determined time, explicitly omitting expectations about technological advancements or changes in future prices of either electricity or appliances themselves. Hausman specifies a static utility-maximization problem that includes both the up-front and operating costs of room air conditioning. He uses data on air conditioner prices, energy efficiency and usage to measure how consumers trade off up-front capital investment with future operating expenses. Comparing the coefficients on unit price and operating costs offers a measure of consumers' willingness to substitute between up-front cost and energy efficiency (higher levels of which lower operating costs in his model). He uses a small but detailed dataset that includes meter data on appliance usage, appliance prices and characteristics, and household characteristics. He concludes that consumers are in fact willing to incur higher purchase costs to save on future operation, but that the implied discount rate (approximately 20 percent) is significantly higher than might be expected.

Dubin and McFadden (1984) allow for endogenous appliance choice each period by assuming conditions that are akin to perfect rental markets. Consumers in their model make concurrent decisions about appliance choice and usage, and can costlessly alter the composition of their appliance stock every period. Their approach broke new ground by allowing key variables to affect both appliance choice and intensity of use. They achieve this by creating an economist-friendly version of the engineering thermal load models, presented in Dubin and McFadden (1983), to make use of variables most often found in household survey data.

Recent models of the durable good choice fall into two categories. The first class reduces the durable good choice into a static optimization problem, in the sense that the modeled relationship between durable good choice and derived energy consumption are unrelated to the role of uncertainty and expectations-formation on the part of the consumer. Bento et al. (2009) and Busse et al. (2009) examine the relationship between demand for cars and demand for gasoline, and each reduces the durable good decision to a static problem in different ways. In Bento et al., the market for cars is assumed to operate as a perfect rental market in which households face no costs of changing their automobile holdings each year. Busse et al. examine demand for cars (and their characteristics) in a reduced form context that abstracts away from the discrete choice facing a given household. They use a transaction-level dataset of car purchases in the U.S. to examine how large swings in gas prices from 2000-2008 affect new and used car purchases, and demand for fuel economy. In this context they explore (among other things) the implied consumer discount rate. Their conclusion is that consumers not only react to changes in fuel costs (by substituting away from gas-guzzling cars), but they may actually over-react. In their study, the implicit discount rate is estimated to be 1.1 percent. All papers (to this author's knowledge) examining the choice of energy-consuming durables use methods that model consumers as static optimizers.

The second class of durable good models explicitly builds uncertainty and rational expectations into a dynamic framework, but until this paper has not been extended to the topic of energy demand. This class of models draws (directly or indirectly) on Rust (1987), which proposes conditions under which the value function that corresponds to (dynamic) discrete choices is the direct solution to a fixed point algorithm. Of these approaches, the framework most similar to the one used in this paper is Prince's (2008) analysis of repeat purchases of personal computers (PCs). Prince uses observed demographic data and information about household stock holdings of PCs to identify drivers of purchase behavior from cross-sectional data. The model also accounts for forward-looking behavior. He finds that price and technology elasticities of demand are mis-measured by 25 to 35 percent when forward-looking behavior is ignored. Other applications of dynamic discrete choice models for durable/storable goods demonstrate the importance of accounting for these dynamic considerations, and include Melnikov (2001) on computer printers, Erdem et al. (2005) on ketchup, Gowrisankaran and Rysman (2009) on camcorders, Schiraldi (2011) on automobiles, and Hendel and Nevo (2006) on laundry detergent (this list is far from exhaustive). Note that many of the methods described above require numerical simulation to allow for unobserved characteristics and/or consumer heterogeneity. My model does not. Because my data include rich variation in household attributes, all heterogeneity in my model is observed.

The drawbacks to building complex structural models of dynamic behavior are clear: such models very quickly become computationally expensive, and often have difficulty accommodating the full array of data available to the researcher. They also often rely on stronger identifying assumptions than their reduced-form or static counterparts. Yet the importance of dynamic considerations (on both the supply and demand sides) in this economic setting is clear, and remains largely ignored in the literature. This paper seeks to begin filling that gap.

3 Data

The primary source of data for this study is the Residential Energy Consumption Survey (RECS), which is conducted every three or four years by the Energy Information Administration (EIA), a division of the U.S. Department of Energy. The sample includes five years of repeated cross-sectional data spanning 1990 to 2005.[5] The RECS data are rich in details about each household's appliance holdings, usage, and other characteristics. I augment the household data with market data on air conditioner prices and efficiency, which I obtain from industry associations and the Census Bureau's Current Industrial Reports.

I use a sub-sample of the RECS data that is intended to isolate appliance choice decision on the part of households. I limit my sample to single-family, owner-occupied houses built before 1987. This eliminates all new construction and allows me to avoid having to disentangle choices about structural housing characteristics related to cooling and energy use (e.g. investment in insulation) from the AC platform choice. I also eliminate households whose owners recently moved.[6] Tables (1) and (3) summarize the RECS household data for the active sub-sample. There are significant differences in characteristics across geographic regions and between households with different air conditioning platforms. Ultimately, my preferred specifications use the entire US with the exception of California.[7] For this reason, I display statistics for this segment in the "US Non-CA" column of Table (1).

In this paper, there are three platform types: central, room, and no AC.[8] Households with central

[5] I use 1990, 1993, 1997, 2001, and 2005.

[6] Similar exclusions are chosen by others using the RECS data, for example, Brill et al. (1999). One might imagine that certain household considerations may be related to the timing of air conditioner purchases in a way that is inconsistent with the model of derived electricity demand that I develop here. For example, central air conditioning is standard in much new housing construction, and tenants often pay for electricity used for production with durables that are purchased by their landlord. The former confounds the determinants of appliance choice with other household decisions, and the latter de-couples the cost of appliance usage from the purchase decision.

[7] California is consistently an outlier on energy policy. To avoid concerns that its inclusion would contaminate the analysis with systematic unobserved heterogeneity, I eliminate it from my preferred subsample.

[8] "Room" units refer to window and wall units, removable or permanent. "Central" units include unitary and split cooling systems. Unitary systems combine the condenser and air handling capabilities in a single casing, whereas

AC tend to have a higher income, a larger home, live in a warmer climate, and have a lower electricity price than houses without central air. One of the most important determinants of air conditioning demand is climate warmth, which here is captured by the variable cooling degree days (C).[9] In general, households in warmer climates (i.e., higher number of cooling degree days) are more likely to own any air conditioner, and are more likely to have central AC.[10] Due to these factors, it is not surprising that households with central AC spend 2.5 times as much per year on cooling electricity as their room-unit counterparts.

An appealing feature of the RECS data is household-specific electricity prices and usage. The EIA obtains actual utility bill data for each household, from which we know (among other things) the total kilowatt hours (kWh) of electricity demand and the dollars spent on electricity. These, in turn, yield an average electricity price. Further, the EIA provides estimates of the electricity used for various household services, including air conditioning.

The electricity price reported in RECS is the average price per kilowatt hour, calculated by dividing electricity expenditures by the usage reported on household electricity bills. There is some measurement error in this variable, both because total expenditure includes some taxes and fees, and in some cases due to the existence of a block-rate structure. Where a block-rate structure exists, marginal electricity price is variable in the sense that increased usage may push a household onto a new, higher segment of the price curve.[11] If consumers know where the kink points are in their price schedule, how much electricity they use each month, and when the utility will visit to read the meter, this is a potential source of endogeneity for the econometrician when trying to find the causal effect of price on usage.

This problem is not new, and generally econometricians select between two approaches. Either they limit the geographic scope of the study in order to obtain actual block-rate schedules (as in Reiss and White (2005)), or they assume that to the that extent consumers respond to price, they do so to average price, not marginal. This is the assumption made, for example, by Dubin

split systems has separate indoor evaporator and outdoor condenser units.

[9]Cooling degree days (C) is a common measure of climate warmth, where $C = \sum_{t=1}^{365} \max\{0, (high_t - 65)\}$, an annual measure equal to the maximum of zero and the sum over each day of the difference between the daily high temperature, $high_i$, in degrees fahrenheit and, in the case of the data used here, 65.

[10]In the RECS data, C is obtained for each household from its local weather station. There are thousands of such weather stations across the country. To disguise the identity of individual households, the EIA adds a random error to each observation of C. This may lead to attenuation bias in the empirical results that follow.

[11]A block-rate structure exists when price is a stepwise function of kWh of electricity usage. For residential electricity, these functions are increasing in usage.

and McFadden (1984), and is consistent with consumers having limited attention, or facing high information acquisition costs. This rationale is explored and supported by Shin (1985). Ito (2010) lends empirical support for the use of average prices in this setting.

Dynamic content of the dataset comes from responses about the age of households' air conditioners. For each household, the data report whether a new AC unit has been purchased within the previous two-year period. That is, if the AC unit is less than two years old, it is considered to be "new". The data do not report the AC model, nor are respondents asked the age of the AC platform that they replaced. To fully describe the state space at the time of purchase, I need to model the household's AC characteristics before the purchase (what I refer to as the "originating state").

I predict the originating state by using aggregate data on household transitions, conditional on pre-existing characteristics. While I cannot identify for a given household whether and what AC it had before a purchase, the data reveal the aggregate fraction of households in each state. Further, households that purchased an air conditioner previously had a finite number of potential originating states from which they made the decision to purchase. Households purchasing a new central unit were either replacing an old unit, in which case their originating state was characterized by central AC, or switching to a central AC platform from room AC. I use estimates of these transition probabilities to match the data in the structural model optimization.[12] I discuss the role of transition paths in detail in Section 4, which describes the model.

I augment the RECS data with information about the prices and efficiency of room and central air conditioning units. The RECS data include information only about the age and type (room versus central) of air conditioner, but nothing to identify the model type or efficiency. As a result, I use a vintage approach to assigning air conditioner attributes. Using the age of the air conditioner in each household, I assign the sales-weighted average efficiency of units sold during that period to the household.[13] As such, I require data on annual average energy efficiency, price, and quantity sold.

Room unit prices and efficiency are available from the Association of Home Appliance Manufacturers (AHAM). These data include U.S. sales-weighted retail prices and efficiency over the period 1987-2005. Figure (1) shows the path of these variables over time. Of particular note is the increase in energy efficiency accompanied by a decrease in the price per unit. It is standard for

[12]See the Appendix for details.

[13]An ideal dataset would include richer information on household-specific AC attributes. The vintage approach a logical alternative in the absence of such data, though it admittedly induces measurement error during estimation.

energy efficiency to be measured differently for room versus central units. Room units are assigned an Energy Efficiency Rating (EER) and central units are assigned a Seasonal Energy Efficiency Rating (SEER).[14]

The data for central air conditioners come from two sources. The U.S. Census provides sales-weighted manufacturing prices and the Air Conditioning and Refrigeration Institute (ARI) provides energy efficiency. The Census publishes central AC prices in their Current Industrial Reports. While ideally retail prices for these units would be available, those data do not exist. Central units are generally not sold through retail channels; rather, they are often included as part of a contractor package with installation and any required dwelling unit retrofitting. Figure (2) shows the evolution of price and energy efficiency for central units of various sizes.[15] As in Hausman (1979), I assume that size (BTU capacity) is exogenously determined by households, primarily by household square footage.

Crucial to the need for a forward-looking demand specification is the fact that air conditioner prices have exhibited large decreases over a period where efficiency has been increasing. A fully informed, rational consumer would incorporate expectations about these trends into his or her decision about if and when to purchase an air conditioner. I proceed under this neoclassical paradigm.[16]

A noteworthy aspect of Figure (2) is the jump in central unit efficiency in 1992, which provides much of the exogenous variation in efficiency that identifies model parameters. The timing of this increase corresponds to the implementation of federal energy efficiency standards that were passed into law as the National Appliance Energy Conservation Act of 1987. It mandated a minimum energy efficiency requirement for twelve types of household appliances to be implemented in 1992. In the five years between the law passing and coming into effect, efficiency of room units increased to the level of the standard so that by 1992 the regulation was non-binding. Growth in central unit SEER was not as rapid, forcing a larger jump in 1992. This event provides exogenous supply-side variation helps to identify demand for efficiency. The absence of a concurrent price jump for central AC suggests that the efficiency increase was not accompanied by proportional cost increases.[17]

[14] SEER equals the ratio of cooling output in BTU to the power consumption in Watts/hr, evaluated over an entire season. EER is a cooling rate equal to BTUh/Watts, evaluated at an outdoor temperature of 95 degrees.

[15] The categories C1-C8 correspond, respectively, to units with kBTU/hr of under 33, 33-39, 39-44, 44-54, 54-65, 65-135, 135-185, 185-250.

[16] One may quite reasonably question whether the neoclassical paradigm is appropriate here, but I do not model alternatives in this paper. See the Results section for a discussion of some results that might increase one's desire to expand the breadth of potential behavioral models.

[17] Engineers cannot point to a single technological development that could account for these changes. A common hypothesis among them, however, is that the 1992 efficiency standards caused manufacturers to focus on improving

Figure (3) shows that the fraction of households with central AC increased significantly in the period around and after 1992. This is suggestive of consumer appetite for energy efficiency, particularly with respect to central AC. The model that follows seeks to explore this relationship.

4 Model

4.1 Basic Setup

I estimate a discrete-time, infinite-horizon dynamic consumer optimization problem. Embedded in this discrete choice framework is a continuous decision about household production, where the durable good can be combined with other inputs to generate a flow of utility. Consumers each period face the choice of whether to purchase (at most) one unit of a durable good, and the quantity of household production. Conditional on the installed durables, households determine their level of production, which serves as an input to solving their dynamic discrete choice problem. If a household makes a purchase, it's stock of the durable is augmented according to the characteristics of the good. One element of the household choice set is the option to do nothing, allowing them to wait for a future period to purchase. Since there is uncertainty about future characteristics of the durables for sale in the market, households operate under uncertainty and may elect to defer their purchase until technology has progressed.

Each durable good is characterized by three state variables: platform type (Z_{it}), energy efficiency (E_{it}), and prices of a room or central AC units (P_t^r and P_t^c respectively). Platforms available to consumers include non-adoption, three levels of room air conditioners corresponding to their number of units, and central air conditioning ($Z_{it} \in \{0, rm_1, rm_2, rm_3, ctr\}$). These states are mutually exclusive.[18] The productive potential of durables persists across periods, but depreciates deterministically. There are no secondary markets, and no transaction costs of purchasing in the

energy efficiency, and that these efforts combined to reduce cost while also achieving compliance under the regulation. For example, adding wrinkles to the heat exchanger material increased efficiency while also allowing manufacturers to use thinner metal. The thinner metal in turn improved efficiency, and cost less. There were also many motor improvements, such as going to a two-speed system (which yields large efficiency increases). Fan blade design also improved. This allowed more air to move with less power, which in turn meant that the fan could be made smaller (another cost improvement).

[18] In the data, some households have both room and central air conditioners concurrently. However, only 2.3% of households in the RECS sample fall into this category. These households have very low average usage of their room AC units (1.2 on a scale of 0 to 4, where 1 indicates "turn on only a few days or nights each year"), so I model them as though they own and use central AC exclusively.

primary market (beyond, of course, the unit price itself). In this model, there is free disposal and the scrap value of the durable is zero.

Household i in time t can be thought of as residing at a point in a multi-dimensional state space defined by their household characteristics (including characteristics of installed durables), and the characteristics of products available for purchase in the market. Market characteristics are denoted by Ω_{it}; housing characteristics by X_i; and endogenous state variables, which characterize their air conditioning appliance stock, by Z_{it} and E_{it}. Market characteristics, Ω_{it}, include air conditioner unit prices and energy efficiency at time t. I assume that elements in X_i are exogenous and constant: electricity prices, income, household square footage and cooling degree days.

Each period in the model corresponds to a duration of two years. At the beginning of each period, consumers have an installed stock of the durable good (as determined by decisions in previous periods, and depreciation), and exhibit other household characteristics. Since I observe these variables for a single year (the final year of the two-year period), there is an implicit assumption that the value of these variables is roughly constant within the period. Consumers face the choice of whether to buy one unit of the durable good. They have perfect information about product characteristics, and may alternately choose to consume a perfectly divisible numeraire good (money).

4.2 Household Production

Households combine electricity with their air conditioners to produce cooling. Households anticipate their derived demand for cooling electricity under each potential AC purchase choice, then select the utility-maximizing option. Utility, u_{it}, is a function of cooling utility, δ_{it}, cooling-related electricity expenditures, $p^e_{it}k_{it}$, AC unit purchase costs, P^j_t (where $j \in \{r, c\}$), and an idiosyncratic shock, ε_{it}. The choice variables here are k_{it}, derived electricity demand for cooling production, and a_{it}, the action embodied by the discrete air conditioner purchase decision.

$$u_{it}(k_{it}, a_{it}) = f(\delta_{it}(k_{it}, a_{it}), p^e_{it}k_{it}, P^j_t, \varepsilon(a_{it})) \tag{1}$$

where

$$k_{it} = k_{it}(X_i, Z_{it}, E_{it}|a_{it})$$
$$a_{it} = a_{it}(X_i, Z_{it}, E_{it}, \Omega_{it})$$

The kilowatt hours of electricity used for cooling, k_{it}, is closely linked to both intensity-of-use and the energy efficiency of the installed air conditioning platform. In each period consumers may choose to alter efficiency by purchasing a new unit. The consumer's action, a, may include the choice of making no purchase ($a = 0$), purchasing a room unit ($a = 1$), or purchasing a central unit ($a = 2$). Specifically, a comes from a choice set, $a \in A(Z_{it})$, where combinations of the available subcollection of $A(Z_{it}) \in \{0, 1, 2\}$ differs based on the installed air conditioner platform, Z_{it}. These differences conform to the restrictions placed on household transitions and imply that not all actions are part of $A(Z_{it})$ for every household. Those with central AC, for example, are restricted from purchasing a room unit, which yields $A(c) = \{0, 2\}$. The option not to buy, $a = 0$, is available to all households at all times.

One component of flow utility, δ_{it}, is derived primarily from climatic comfort, T_{it}, which can be thought of as a negative function of the temperature difference between the achieved temperature and the household's conceptual "bliss point". Flow utility is also derived from unobservable attributes of the air conditioner platforms themselves, factors which may include humidity reduction, noise, aesthetics, seasonal installation and storage costs, and maintenance expenses. In a linear specification, these could enter δ_{it} as platform dummies (D^r and D^c), though they could enter in a more flexible way.

Over long time periods, cooling degree days, C_i, reflects the "hotness" of climate relative to 65 degrees.[19] I assume that it also equals the desired level of cooling that a household living in that climate would want to produce. Where $\widehat{C}_{it}$ is the amount of cooling produced by household i, I define T_{it} as

$$T_{it}(C_i, k_{it}) = -(C_i - \widehat{C_{it}}(k_{it})) \tag{2}$$

with

$$\widehat{C_{it}}(k_{it}) = \gamma^j k_{it} E_{it} / S_i. \tag{3}$$

The effectiveness with which this energy is converted into cooling is increasing in E_{it}, the energy efficiency of the AC unit, and γ^j, the platform-specific factor, S_{it}, and decreasing in the square footage of living space being cooled. The specification in (3) assumes that cooling output (in BTUs) is converted into degree-days of cooling at a rate inversely proportional to the square footage of the area being cooled.

[19]Recall that cooling degree days (C) is calculated according to $C = \sum_{t=1}^{365} \max\{0, (high_t - 65)\}$.

The factor γ^j converts $k_{it}E_{it}/S_{it}$ from a measure of BTUs of cooling production per square foot of cooling area into cooling degree days. It is specific to the AC platform being used by household i, where, again, $j \in \{r, c\}$ corresponds to room or central. Household characteristics such as quality of insulation, humidity, heat transfer capability, number of windows, and use of space may cause there to be differences in γ^j between households. Further, differences may exist between air conditioning platform types. There are at least two reasons to believe that $\gamma^c < \gamma^r$. Cooling energy is lost when transmitted via ducts, which are never perfectly insulated. Additionally, central AC concurrently cools multiple rooms in a home, regardless of their occupancy. Multiple room units distributed across these rooms would result in less wasted cooling if used only when needed. That is, the first best level of cooling can only be achieved via use of room units unless every room in the structure is occupied.[20] In the base empirical specification I set $\gamma^c = 0.7\gamma^r$, and later show that my results are robust to changes in this assumption.

In combination, equations (1)-(3) imply a framework for deriving k_{it}, the conditional demand for (cooling) electricity. The choice variable is k_{it}, the kilowatt hours of electricity used,

$$k_{it} = \arg\max_{k} \; u_{it}(k_{it}|a_{it}), \tag{4}$$

where it is recognized that the optimal choice of k_{it} is conditional on the action, a_{it}, chosen by the household.

In the results that follow, I allow for a flexible functional form for the derived demand equation and allow for a broad interpretation of the relationship between k_{it} and each of the relevant state variables. The first-stage regression takes a linear form, relating k_{it} to a second-order polynomial function, Φ^j of the state variables separately for each type of air conditioning platform.

$$k^r_{it} = b^r_0 + b^r_E E^r_{it} + b^r_\Phi \Phi^r_{it} + \epsilon^r_{it} \tag{5}$$

$$k^c_{it} = b^c_0 + b^c_E E^c_{it} + b^c_\Phi \Phi^c_{it} + \epsilon^c_{it} \tag{6}$$

I perform this regression separately for households with room units and central. The arguments of functions Φ^j include household income, square footage, CDD, p^e, and in the case of households with room air conditioners, dummy variables for the number of units. The parameter estimates from these regressions are used to predict $\widehat{k_{it}}$, which vary with household characteristics and with the installed air conditioner platform. Using the estimate $\widehat{k_{it}}$ one can calculate $\widehat{C_{it}}$ and, in turn,

[20] I thank Mark Modera, Director of the UC Davis Western Cooling Efficiency Center, for these observations.

T_{it}.

4.3 Durable Good Purchase Decision

I assume that utility is log-linear in cooling utility, δ_{it}, and AC-related expenditures, $k_{it}p^e_{it}$ and P^j_t. This yields the following functional forms:

$$u_{it}(\hat{k_{it}}|a_{it}) = \left\{ \begin{array}{ll} \delta_{it}(\widehat{k_{it}}) + \alpha_1 \ln(p^e_{it}\widehat{k_{it}}) + \varepsilon_{it} & \text{if } a_{it} = 0 \\ \delta_{it}(\widehat{k_{it}}) + \alpha_1 \ln(p^e_{it}\widehat{k_{it}} + P^r_t) + \varepsilon_{it} & \text{if } a_{it} = 1 \\ \delta_{it}(\widehat{k_{it}}) + \alpha_1 \ln(p^e_{it}\widehat{k_{it}} + P^c_t) + \varepsilon_{it} & \text{if } a_{it} = 2 \end{array} \right\} \tag{7}$$

and

$$\delta_{it} = \alpha_2 \ln T_{it}(\widehat{k_{it}}|a_{it}) + \alpha_3 D^r_{it} + \alpha_4 D^c_{it} \tag{8}$$

Equations (7) and (8) imply a theoretical relationship between the state variables and k_{it}, but this relationship does not have an adequately-simple reduced form. However, the flexible first-stage regression described above is intended to approximate the essential aspects of the relationship.

I assume that the utility from cooling of those with no AC (the "outside option" in this model) is equal to $(-\alpha_2 \ln C_i)$. As such, households with no air conditioning experience lower utility in warmer climates. I further normalize the flow utility of the outside option to zero, which is standard in the literature and necessary for convergence of the maximum likelihood estimator. In the specification outlined in (7) and (8) above, there are four structural parameters to estimate: α_1, α_2, α_3, and α_4.

Upgrading one's air conditioning holdings changes derived electricity demand differently for households with room units versus central. In the case of room units, purchasing an additional unit affects electricity use through increased energy efficiency and through the presence of an additional AC unit. Electricity use for central AC changes only through the difference in energy efficiency.

When a household considers purchasing a new appliance, it implicitly predicts its well-being under each scenario: buying and waiting. This creates an empirical challenge since each household's electricity usage, k_{it}, is observed under only one of the potential decision scenarios for each household. To remedy the lack of complete data, $\widehat{k_{it}}$ must be estimated under the unobserved alternatives, which I do using the results of the household production optimization above (which predicts optimal cooling and electricity usage conditional on household attributes).

One challenge inherent to modeling household air conditioning demand is capturing the essential features of transition paths between platforms while still maintaining computational feasibility. I do this by constraining the potential transitions available to households in the model. Those starting with no AC first must visit the room AC state before entering the central AC state. Also, only one room unit can be purchased in a single period. The maximum number of room units allowed in my model is three.[21] Households that have already adopted central AC have the choice of replacing their current unit with a new one or doing nothing. These restrictions make the model tractable while still allowing for a rich set of stock evolution patterns.

The units for $E_{it} \in \{E^r_{it}, E^c_{it}\}$ are the energy efficiency ratio (EER) for room units and the seasonal energy efficiency ratio (SEER) for central AC respectively. Each of these is a measure of the efficiency with which electricity entering the units (in kWh) is converted into cooling energy (British Thermal Units or BTUs).

I assume that all units are purchased new, are infinitely-lived, and that there are no maintenance expenses. Further, I assume that AC units experience an annual efficiency deterioration of $\rho = 0.01$.[22] The advertised energy efficiency of a given air conditioner is measured under optimal operating conditions while brand new. Over time, even if optimal external conditions are achieved (which is rare), efficiency deteriorates.[23] This deterioration implies that older units use more electricity to generate a given amount of cooling than new units. This holds even if there were no technological advancement improving energy efficiency over time. Scrappage is assumed to be costless and in this model there is no distinction between allowing a unit to lay idle or physically disposing of it.

Consumers determine whether or not to purchase a new air conditioner according to a forward-looking value function that I represent in Bellman form:

$$V(\varepsilon(Z_{it}), X_i, Z_{it}, E_{it}\Omega_{it}) = \max_{a_t \in A(Z_{it})} \{u(a_t, k_{it}) + \beta E[V(\varepsilon_{i,t+1}(a_t), X_i, Z_{i,t+1}, E_{i,t+1}(a_t), \Omega_{i,t+1})|\Omega_{i,t}]\} \tag{9}$$

[21]Of households in the data, 2.3 percent have more than three room units.

[22]I thank HVAC engineers Jim Crawford (Trane) and Donald Karner for their assistance with reaching this figure, and for their patience while explaining the intricacies of heat transfer to me.

[23]The cooling fluid that is used and re-used in the cooling cycle may leak, the heat transfer surface (which mobilizes the cooling created by evaporating the cooling agent) can become dirty over time, and the compressor pistons wear. All of these may lead to reduced cooling efficiency.

where β is the consumer discount factor (assumed to be 0.9 in the estimation) and Ω is assumed to follow a first order Markov process. As observed by Rust (1987), the solution to this value function requires integration over the realizations of two random variables: $\varepsilon_{it}(= \varepsilon(Z_{it}))$ and Ω_{it}. Even if we were to discretize over the support of each and calculate the value function for each bin, dimensionality becomes a severe problem. Rust assumes "conditional independence" between these two variables so that the solution can be separated into the calculation of two distinct integrals, one that can be solved as a fixed point and the other as a numerical optimization after discretizing the state space. In the context of this model, conditional independence is satisfied implicitly by assuming that draws, ε_{it}, from the Extreme Value Type I distribution are *iid*. It implies that all information about the shock in period $(t-1)$ is transmitted into period (t) via the household's action.

One can now separately consider the expectation of the value function after having been integrated over realizations of the unobservables, ε_{it}. Let this term be defined as

$$EV_i = \int_{\varepsilon_i} V_i(\varepsilon_i, X_i, \Omega_i, Z_i, E_i|\alpha)dF_{\varepsilon}. \tag{10}$$

Expressing EV_i in this way allows it to be computed as a fixed point of a separate contraction mapping, as shown in (11). I solve for the fixed point using successive approximations, with convergence assumed when the maximum difference falls below the threshold 10^{-9}.[24]

$$EV_i = \ln\left(\sum_{a\in A(Z_i,E_i)} \exp(u(a,k_i) + \beta E[EV_i(a)|X_i, \Omega_i, Z_i, E_i, \alpha])\right). \tag{11}$$

Having solved for the value function, EV_i, the postulated distribution of ε_i enables one to write the probability of each action in a household's choice set as a standard logit probability:

$$\widehat{p}_{it}^{\widetilde{a}}(\alpha) \equiv \widehat{p}_{it}^{\widetilde{a}}(EV_i) = \frac{\exp(u(\widetilde{a},k_i) + \beta E[EV_i(\widetilde{a})|X_i, \Omega_{it}, Z_{it}, \alpha])}{\sum\limits_{a\in A(Z_{it})} \exp(u(a,k_i) + \beta E[EV_i(a)|X_i, \Omega_{it}, Z_{it}, \alpha])}. \tag{12}$$

Since all $a \in A(Z_{it})$ are mutually exclusive and exhaustive, the probability of a household choosing

[24]See Rust (1987) for a detailed description of the fixed point algorithm.

not to buy is

$$\widehat{p}_{it}^{0}(\alpha) = 1 - \sum_{a \in A(Z_{it})^{-0}} \widehat{p}_{it}^{a}(\alpha).$$

The probabilities calculated in this way are associated with stylized households whose characteristics are a result of a variable discretization. Following the methodology outlined in Rust, I discretize the state space into a finite number of bins, the values in which span the range of the variables in my sample and any forecast range that is being considered. This is also consistent with Hausman (1979), who finds that his results are insensitive to the limits used to define the bins.[25]

I discretize cooling degree days and electricity price into four bins, energy efficiency (separately for new units and existing), AC purchase prices, and house square footage into three, and income into two. There are thus 96 different household "types" in static variable space (4x4x3x2 for CDD, electricity price, square footage, and income, respectively). The dimension of each household's dynamic state space depends on it's endowment of air conditioner platform. Households with no AC have 9 potential price and efficiency states that correspond to a purchase of room AC (3x3 for efficiency and price of room unit). Households with one or two room ACs reside in one of the existing-AC efficiency bins, and can purchase either another room AC or a central unit of a given price and efficiency. Thus there are 108 (= 2*3*(3*3 + 3*3)) possible stylized households of this kind. Households with three room units are constrained to purchase only central AC, and thus can reside in one of 27 bins. The state space of households with central AC also has 27 bins. The total state space thus has a dimension of 16,416 (= 96*(9+108+27+27)). I calculate the value function associated with each of these stylized households.

To generate a continuous likelihood function, I use linear interpolation to construct a continuous value function for each choice available to the households, and then construct the probability of each choice as in (12). Linear interpolation provides smoothness for the likelihood function as well as helping to exploit heterogeneity in my data. When, for example, the number of bins into which I discretize a given variable is small, interpolation allows for assignment of different probabilities to households whose characteristics are different but fall into the same bins.

[25] Hausman (1979) groups air conditioners into three bins: high-efficiency, medium-efficiency, and low-efficiency.

4.4 Markov Transition Probabilities

There are four dynamic state variables that consumers in the model forecast in order to determine whether to buy today or wait: energy efficiency and price for both room and central units. I initially included electricity price as a dynamic state variable, but the computational cost of the larger state space outweighed the benefit: while electricity prices fluctuate over time, they do not exhibit a consistent long-run trend over the period examined here. Rational expectations over electricity prices, therefore, do not differ greatly from the naïve expectation that tomorrow's price will equal today's.[26]

In the model, air conditioner unit efficiency and price are market-wide variables common to all consumers, and changes in each will influence behavior by altering the trade-off between up-front costs and ongoing operating expense. Increased energy efficiency of one's AC stock is expected to affect intensity of use in the same direction as a decrease in p_i^e. The paths of these market variables are different for room and central units. Each of these variables evolves over time as a result of market forces and/or technological change, and households formulate predictions about the nature of this evolution. I assume that each of these four dynamic variables follows a first-order Markov process.[27]

Recall that Ω_{it} represents the space spanned by the four dynamic market variables. The subscript i is required to indicate that Ω_{it} is household-specific. Different-sized houses face different prices for central AC units. One of the major factors that determines the price of central AC is how much cooling power (in BTUs) the unit has. All else equal, households with large square footage of cooling space will require a larger central AC unit and thus face higher prices in the marketplace. Allowing for idiosyncratic Ω_{it}'s is important for maintaining the relevance of central AC prices in my model. The range of these prices over all home sizes is extremely large, forcing the transition matrix for any reasonable number of discretized variable bins to be the identity matrix. However, since the prices vary around nodes that are related to housing square footage, idiosyncratic transitions on this dimension capture the information that households can reasonably be expected to use.

The variables being forecasted by consumers here are modeled to have a simple dynamic structure.

[26]The estimated transition matrix for electricity prices was the identity matrix when the range spanned by p^e was discretized into a reasonable number of bins.

[27]Some studies (originally Melnikov (2001), and later Gowrisankaran and Rysman (2009), Schiraldi (forthcoming), and Hendel and Nevo (2006)) reduce multiple dynamic variables into a (logit) inclusive value, which they then assume follows a first-order Markov process.

I specify a first-order autoregressive process for each dynamic variable v in Ω:

$$\log(v_t) = a_0 + a_1 \log(v_{t-1}) + e_t, \qquad e_t \sim N(0, \sigma_v^2) \tag{13}$$

One might imagine putting other variables besides v_{t-1} as explanatory variables, but they appear to be unimportant in specification testing. This facilitates the computation of the transition matrix for Ω_{it} by allowing the Markov transition probabilities of each variable to be calculated independently. Using the first and second moments of the predicted errors from (13), $\widehat{e_t}$, transition probabilities can be calculated for the range spanned by each bin of the dynamic variable, as in Tauchen (1986). The resulting transition matrices replace the expectation in equation (11) in discretized state space.

4.5 Identification

4.5.1 Identifying Unit Price and Efficiency Effects

Separate identification of the effects price and efficiency on demand for AC unites requires that these two variables exhibit independent variation. A quick glance at Figures (1) and (2) shows that there is clear negative correlation, but a closer inspection reveals significant variation in price that remains after having controlled for efficiency (and vice versa). The sources of this residual variation are twofold. First, energy efficiency (the rate at which energy is converted to cooling) is unrelated to cooling capacity (which is a monotonic function of the level of energy used as an input), and prices increase in cooling capacity. Therefore, for a given unit efficiency, different size units will have different prices. Further, residents of larger homes require larger units. Under the assumption that AC cooling capacity is exogenous to the household, any cross-sectional variation in the size requirement is also independent with respect to energy efficiency. The cross-sectional variation in prices across unit size is seen clearly in Figure (2).

The second source of independent variation between price and efficiency comes from the fact that, even after conditioning on unit size, there is independent variation remaining between price and efficiency. The correlation coefficient between price and efficiency within a size class ranges from -0.53 to -0.88, and has an average of -0.79 across the eight central unit sizes. Residual variation (absent only if these statistics are +1 or -1), combined with the assumption that efficiency and price of AC units are taken to be exogenous with respect to the consumer (which is realistic, given the use of micro-data), allows for separate identification of the price and efficiency effects.

4.5.2 Identifying Derived Demand for Cooling

Identification requirements for demand on the intensive margin are standard. Consistent estimation of equations (5) and (6) requires that ϵ is orthogonal to the other right-hand side variables. The efficiency variables, E^r and E^c, are most likely to violate this assumption, since households with unobservable characteristics that lead to a high appetite for cooling services are also likely to purchase more efficient units. The existence of a rebound effect is also a concern when trying to reach a consistent estimate of b_E. [28] Ideally, efficiency of installed air conditioners would be randomly assigned across households, allowing for the rebound effect to be estimated directly. Of course, efficiency is a choice variable of interest in the dynamic model, so no valid instrumental variable exists.

To overcome both problems (endogeneity and the rebound effect), I allow efficiency to enter these first-stage regressions only as a linear term, then exploit the fact that the coefficient on the linear efficiency term will equal $-(1 - rb)$, where rb is the true rebound effect. I use estimates of the rebound effect from the literature to restrict the coefficients b_E^r and b_E^c. The most directly relevant estimate of the rebound effect for air conditioner usage is by Dubin et al. (1986), who use data from a controlled experiment by a Florida utility. They estimate it to be between 2 and 13 percent depending on the month of the year. I use the midpoint of these, 7.5 percent, which implies that, all else equal, a one percent increase in energy efficiency will lead to a 0.925 percent decrease in electricity use. Therefore, for my preferred specification I restrict the coefficient on efficiency to -0.925, and later show that my results are robust to changes in this assumption.

Restricting the coefficient on efficiency while allowing the estimated response to electricity price has implications for the level of rationality assumed by the model. A fully rational, cost-minimizing household would respond identically (in terms of demand for cooling) to a given increase in energy efficiency and a equivalent decrease in the electricity price. To see this, recall again that electricity is simply an input to cooling production, and that energy efficiency is a technology that makes production more efficient. However, there are compelling reasons believe that the rational, fully informed consumer may not be the appropriate starting point. For example, if there is asymmetric salience of efficiency and price to the consumer (due to costly information acquisition, perhaps), one would not expect to observe a symmetric response to changes in these variables.

[28] Recall that increasing efficiency reduces the cost of cooling production. Moving along the demand curve will result in a higher quantity of cooling.

In any case, even after having accounted for potential bias to the coefficient on efficiency, the orthogonality assumption, $E[X_i, Z_i, E_i|\epsilon] = 0$, remains strong. It requires that there are no unobserved household characteristics that correlate electricity usage and observed household characteristics (income, square footage, climate, and electricity price). One can make a legitimate argument that some variables (e.g. quality of insulation or number of windows) about which the RECS provides household-level data belong in the conditional demand specification. Unfortunately, even with the available RECS data, there is a steep tradeoff between computational feasibility offered by model parsimony, and inclusion of all potentially-important variables. As a result, the importance of accounting for expectations and uncertainty in durable good demand analyses must be weighed against these tradeoffs. I proceed under the required assumption of orthogonality.

4.6 Computation

I calculate the probability of the observed behavior of households in my dataset as outlined above and solve for the structural parameters, α, via maximum likelihood.

$$\log L(\alpha) = \sum_{n=1}^{N} w_n \log(\widehat{p}_n^a(\alpha)) \tag{14}$$

Specifically, I maximize the log-likelihood function (14) with respect to α using a quasi-Newton method (BFGS), where w_n is the sample weight corresponding to household n.

Computation of each MLE estimate first uses a compass search, then proceeds to the gradient-based Newton-Raphson method. The compass search is robust to potentially poor starting values, but is far slower to converge than the more sophisticated gradient-based method when in the neighborhood of the solution. It requires between four and five hours of processor time to converge (at a tolerance of e^{-5}). I generate a bootstrapped sample of 200 estimates for each of the six regions, which requires 35-40 days of processor time per region. Bootstrap starting values are parameter estimates from the full sample estimation. Summing over all regions, the processor time required to calculate the estimates presented below is 200-250 days. Bootstrap simulations of the elasticities are significantly faster: each counterfactual scenario projected out five periods (to 2015) requires ten minutes. After bootstrapping (again using a sample of 200), the elasticities presented below took five computation days.

5 Results

Results from the model describe households that are forward-looking, that value the stream of savings derived from investing in energy efficiency, and that are engaged in their choice of air conditioner purchase timing. Conservation by these households is observed on both the intensive and extensive margins. High energy prices lead to immediate reductions in intensity of air conditioner use. Increases in energy efficiency accelerate purchases of durables, and lead to gradual but ever-larger reductions in energy use. These results indicate that both carbon taxes and efficiency standards ought to be effective mechanisms for achieving energy conservation. On the other hand, own-price elasticity of the durables is low, implying that new appliance rebates are likely to provide weak demand response and little corresponding energy reduction. In this section I discuss the results and implications of the derived demand analysis, the structural parameter estimates, summary statistics (elasticities) implied by the model, model fit, policy counterfactual simulations and, finally, consumer welfare implications.

5.1 First-Stage Derived Demand Analysis

The derived demand for cooling electricity is estimated using equations (5) and (6). Cross-sectional variation in price, efficiency, and household attributes pin down the OLS estimates. As discussed earlier, consistent estimates are obtained if unobserved determinants of demand for cooling are orthogonal to the observable characteristics. In an unrestricted specification, this is likely not to be the case for the energy efficiency. I address this concern by restricting the coefficients b_E^r and b_E^c based on estimates from the literature. Under the assumption that this restriction is correct, and that consumers respond to average (not marginal) electricity prices, equations (5) and (6) will yield consistent estimates of derived demand for cooling electricity.

Results from the electricity demand analysis are presented in Tables (4) and (5). Geographically, my preferred specification is "US Non-CA", which includes the lower forty-eight states minus the state of California. For decades, California has been an outlier on energy policy, and its efforts in this area are likely to have induced systematic unobserved heterogeneity. If true, this could bias the results, and focusing on the rest of the country eliminates this concern. By removing California from the national sample, the apparent elasticity of cooling with respect to electricity prices drops by nearly 60 percent. This large difference may be due to increased behavioral response in California resulting from conservation awareness (or similar) programs.[29] In any case, I present the reduced form results for both the entire U.S. and excluding California.

The specification is linear in squares and interaction terms of the state variables, as described in equations (5) and (6). Tables (4) and (5) present the marginal effects of each variable (in logs) on derived electricity demand from cooling. Results fall within a range that is consistent with the literature. The implied (short run) price elasticity of demand for cooling electricity is -0.07 for central AC and -0.34 for room AC. This lies near the range of -0.29 and -0.49 estimated by Reiss and White (2005) for the sensitivity of total household electricity demand to price. Climate warmth (as measured by cooling degree days) is a strong determinant of cooling demand, with an elasticity near one for the nation as a whole. Cooling energy increases in square footage for households with central AC, but not for room AC. Income is associated with higher cooling demand. For households with room air conditioners, adding a second unit is associated with just over a 50 percent increase in cooling electricity. The third unit adds an incremental 20 percent.

Recall that the purpose of the derived demand analysis is to predict usage under unobserved household states, so that each element of the consumer choice set is defined. The direct retransformation (from logs to levels) of predicted values of electricity usage induces a systematic bias that is well-known in the literature. To account for this, I apply a smearing adjustment as described in Manning (1998).

5.2 Structural Model Results and Implications

Estimates of the model's structural parameters are displayed in Table (6). By themselves, these coefficients have little meaning. Their value is derived from using them as inputs into simulations that allow for the computation of economically-relevant statistics, such as elasticities or welfare changes. The signs and significance are nonetheless revealing. The coefficients on comfort and expenditures are positive and negative, respectively, and both statistically significant. The positive room platform dummy coefficient implies that certain attractive unobservable characteristics (i.e., not energy efficiency), such as humidity reduction, outweigh others such as noise or poor aesthetics. Standard errors are bootstrapped to account for estimation error in the first of the two-step procedure. These are calculated by drawing 200 samples with replacement, and solving the maximum

[29] No similar disparity is observed in the marginal effects on Room AC households, but this is likely due to the relatively small number of households in California with room units.

likelihood estimator for each draw.

The coefficient on the central AC platform dummy variable is unidentified. As it only affects households that upgrade from room to central, I thus cannot distinguish upgraders from repeat purchasers in the data. I solve this by calibrating α_4 via constrained maximum likelihood to match a micro moment. This procedure equates the predicted ratio of central replacements to upgrades from room ACs to the observed aggregate ratio in the data.[30] This process is far too computationally-expensive to replicate for each bootstrap. Instead, the bootstraps are computed with this parameter fixed, which allows the predicted moment to deviate from its sample counterpart. The properties of this moment can be seen in Table (10).

5.3 Model Fit

Here I examine two aspects of model fit. First, I compare the unconditional discrete choice probabilities from the data to those predicted by the model. This provides a rough confirmation of the predictive power of the model. I then test the goodness-of-fit of the rational expectations hypothesis against two alternative and less-sophisticated hypotheses of consumer behavior. The benefits of this exercise are two-fold: it aids us in interpreting the structural model and counterfactual results, and provides justification for the dynamic structural approach to demand estimation in this setting.

Table (11) shows the annual unconditional choice probabilities in the data and those predicted by the model. From 1991-1993, the period spanning the year in which the national energy efficiency standards became binding for central air conditioners, the fraction of households purchasing central units was not markedly higher than in other years, as one might expect. During this period, the US was experiencing a recession, which put downward pressure on consumer durable good purchases.[31] Since the model controls for income, the recession effect is internalized in the predicted probabilities. On average, the model slightly over-predicts the purchase of air conditioners on all years but 2005, with the differences mostly attributable to room AC units. Overall, the model is quite good at predicting central AC purchases and no purchases, with the increasing purchase probability over time in the data being reflected in the predicted choice probabilities.

Next, I test the assumption of rational expectations against two alternate behavioral hypotheses:

[30]The process by which the aggregate moments are calculated is described in the Appendix.

[31]According to data from the Philadelphia Federal Reserve Bank, average quarterly expenditures on consumer durables in 1991 and 1992 were 4.1 percent lower than in 1990 and 9.9 percent lower than in 1993.

myopic consumers and naïve expectations. This is done by imposing each alternative behavioral hypotheses on the model, then testing whether rational expectations or the alternatives best fits the model to the data. "Myopic consumers" are defined to have a discount rate of 100 percent, and thus do not place any value on utility beyond the current period. Consumers with "naïve" expectations ignore trends of key dynamic variables, and act according to the belief that efficiency and unit prices follow a random walk. I find that each of these alternative behavioral models is weakly rejected in favor of the null hypothesis, rational expectations.

Table (12) displays the likelihood ratio statistics that compare the fit of these models.[32] The statistics are distributed χ-square with degrees of freedom equal to the number of restrictions imposed on the base model when estimating under the alternative hypothesis. The myopic specification applies a discount factor of $\beta = 0$, which reduces the dynamic model down to a static logit choice. This model is overwhelmingly rejected in favor of the null hypothesis. Strong evidence from the likelihood ratio test is bolstered by coefficient estimates that make no sense (e.g. a positive coefficient on expenditure). The naïve expectations case is mechanically equivalent to setting $a_0 = 0$ and $a_1 = 1$ in Equation (13) for each of the four dynamic variables. In this comparison, rational expectations fits the data slightly better than naïve expectations, and for the US sample, naïve expectations can be rejected with 90 percent confidence.

5.4 Counterfactuals

Parameter estimates from the structural model provide a foundation for analyzing several counterfactual scenarios, such as the change in demand for air conditioners that one might expect if energy efficiency of available units increases, or if price decreases. Results from these counterfactual scenarios are summarized as elasticities that, until now, are in short supply in the literature. The reference case (to which I refer below as the "base case") for the elasticity calculations begins in 2005, immediately after the end of the final period in the data. Households from the 2005 RECS subsample comprise the starting point. The dynamic variables (unit price and efficiency) are projected forward according to the AR1 process in equation (13), which governs the transition of dynamic state variables during estimation. Each household is then "divided" according to its discrete choice probability and sample weight, creating what amounts to new households in

[32]The likelihood ratio test is not exactly correct for this context, in which the models are non-nested. However, it is suggestive of relative model fit.

subsequent periods.[33]

To evaluate the effects of price changes and efficiency standards, a small perturbation is applied to each variable of interest. The counterfactual discrete choice probabilities are calculated with and without the change, and the paths of state variables and household outcomes adjusted and tracked for each future period. For energy efficiency and AC unit price scenarios, I alter the characteristics of room and central units simultaneously. Each household's air conditioner investments and energy use are then projected several periods into the future, as in the base case. The elasticities are calculated by comparing predicted behavior under these scenarios with behavior in the base case.[34]

The change in the perturbed variable is modeled as being permanent and unexpected to the consumer. After the unexpected shock occurs, consumers incorporate the information into their rational expectations of future realizations (via the AR1 process for dynamic state variables, energy efficiency and unit prices, or as a new, permanent level for the static state variable, electricity price). Standard errors on the elasticity estimates must be bootstrapped for analogous reasons to the structural parameter estimates. The bootstrap is implemented using the same draws from the structural model estimation procedure.

Results presented below are divided into per-period and cumulative elasticities. Per-period elasticities are calculated by comparing air conditioner purchase behavior in the given year to behavior observed under the base case in that year. These elasticities will reflect fluctuations due to delay or acceleration of purchases, and offer insights into behavioral responses at higher-frequency. They can be thought of as being somewhat analogous to patterns of short-run storable goods purchases (which, for example, are accelerated during sales).[35] Cumulative elasticities, on the other hand, smooth over purchase accelerations and delays. To calculate these, the elasticity in a given year is calculated with respect to cumulative base case demand from 2007 through the year in question.

[33] With each additional future period, the number of households grows exponentially as a result of this subdivision. This is necessary to keep track of the state variables associated with households along each feasible sequence of discrete choices.

[34] Formally, the per period elasticity in of, say, electricity demand, k_T with respect to an immediate and permanent change in energy efficiency, E, starting in 2007 is equal to

$$\varepsilon_T^{kE} = \frac{\frac{dkWh_T}{kWh_T}}{\frac{dE}{E}}.$$

The cumulative analog is

$$\varepsilon_{Tcum}^{kE} = \frac{\frac{d\sum_{t=2007}^{T} kWh_t}{\sum_{t=2007}^{T} kWh_t}}{\frac{dE}{E}}.$$

[35] See Erdem et al. (2003) and Hendel and Nevo (2006).

All elasticities should be interpreted as market elasticities, as opposed to elasticities for individual products.[36]

5.4.1 Discussion of Counterfactual Results

Table (7) displays cumulative elasticities of demand for air conditioners with respect to energy efficiency, unit price, and electricity price. These cumulative elasticities reflect enduring changes to the composition of the installed AC stock, and are thus more appropriate for considering long-run effects of policy changes on energy demand than their period-by-period analogs shown in Table (8). I focus the discussion that follows on the cumulative results, due to their greater policy-relevance. These statistics are calculated at two-year intervals when simulating demand forward ten years.

The high elasticities of demand with respect to energy efficiency are consistent with households that consider future operating expenses when making their purchase decision. Further, while demand is not entirely inelastic with respect to the price of the durables itself, it is far less elastic than one would expect a myopic consumer to be. Results also indicate that consumers purchase fewer air conditioners when the expected future operating costs are high. Overall, though, demand is less than unit-elastic with respect to all variables. The elasticity of unit demand with respect to efficiency is between 0.7 and 1.0 for central, and from 0.2 to 0.3 for room units. Efficiency increases will translate into growing electricity savings on aggregate over time. The steady increase in the usage elasticity (from -0.1 to -0.3) reflects the slow transmission of conservation that is inherent when efficiency must be purchased in discrete chunks as a characteristic embedded in new capital.

One might expect that efficiency standards (of a fixed magnitude across platforms) would affect demand for room and central units asymmetrically, with energy-intense central platforms becoming relatively cheaper to operate, and thus more attractive. If true, increasing efficiency may induce households to upgrade from room to central. If the differential in electricity usage between central and room units outweighs the energy savings from higher-efficiency units, aggregate electricity usage from cooling would actually *increase* in this scenario. This effect is very similar in nature to the "rebound effect", except that it operates on the extensive margin. Because the model estimated here tracks the AC purchases of individual households, I can test the extent to which this ought to be a concern. I find that, while present, this effect is negligible. This is the case primarily because most of the increased demand due to higher efficiency comes from replacements of existing (less

[36] From the perspective of energy markets and environmental considerations, the market elasticity is most relevant.

efficient) central units, which yields substantial decreases in usage.

Households exhibit a moderate response to AC unit prices, but far less so than to energy efficiency. The own-price elasticity of unit demand is -0.2 and -0.1 for room and central units, respectively. The aggregate impact on electricity demand from changes in AC prices however, (e.g. via appliance rebates) is theoretically ambiguous. To the extent that conservation gains from replacement of old, inefficient units outweigh the increase in electricity demand from first-time adopters or upgrades to more energy-intense platforms, electricity demand will fall. Estimates from my model show the former effect to slightly outweigh the latter, implying that new appliance rebates will achieve net (albeit moderate) energy savings. This effect would be magnified if rebates focused on only energy efficient units (as is common with current EnergyStar rebates), or if scrappage subsidies were offered for the most inefficient installed units. Given the small magnitude of unit prices on long-run electricity demand, though, it is likely that policies based on efficiency or electricity prices would yield more conservation benefits.

While energy efficiency standards and unit prices operate via the extensive margin, electricity prices most strongly influence intensity of use. I quantify the extent to which this intensive margin response is transmitted into extensive margin decisions. Theory does not predict whether demand for AC units should respond positively or negatively to changes in electricity price. When electricity becomes more expensive, there are resulting effects that act in different directions. All households will move along their demand curve and consume less on the intensive margin. On one hand, demand for new, efficient units should increase due to replacement of inefficient units as electricity becomes more expensive. However, overall demand for air conditioning may fall with higher operating prices, and upgrading to more energy-intensive platforms (to central, or even just a higher number of room units) becomes less attractive. My results indicate that the latter effect dominates, and that unit demand decreases in electricity prices (though only slightly, and statistically indistinguishable from zero for central). The intensive margin response dominates that of the extensive margin, and aggregate electricity use changes only slightly more in the long run than in the short run.

These elasticity estimates are robust to changes in assumptions about some of the key input parameters. Table (9) displays estimates of the counterfactual elasticities derived under alternate assumptions about the rebound effect, discount rate, and relative platform efficiency (γ^c/γ^r). Nine elasticity estimates are computed for each of the five alternate specifications. Of these 45 statistics, 43 (95.6 percent) lie within the 95 percent confidence interval of the initial estimates.

5.5 Welfare Implications and Policy

A major advantage of the structural model is the ability to conduct welfare analysis, thereby assigning a monetary equivalent to the utility implications of proposed counterfactual scenarios. I compare the consumer welfare implications of a (consumer-borne) electricity tax to a chance in the energy efficiency of AC units available for purchase. The welfare analysis that I present here involves two steps. First I calculate the energy efficiency change that is (consumer) welfare-equivalent to a five percent electricity tax, assuming a utilitarian planner's objective function. To avoid having to make any unnecessary presuppositions about symmetry of the consumer welfare function, I estimate an energy efficiency change that will operate in the same direction as the electricity tax. I find that a 5.0 percent increase in electricity prices is consumer welfare-equivalent to a decrease in energy efficiency of 31.6 percent. The relative magnitude of the efficiency change is due to the different size of affected cohorts: only future AC purchasers are ultimately affected by the efficiency change, whereas all current AC owners are affected by the electricity price increase.[37]

Next, for both the electricity tax and efficiency change, I calculate the compensating variation – the monetary transfer required to return households to their original base case utility. Table (13) presents the compensating variation for the five percent electricity price increase and 31.7 percent energy efficiency decrease. The median (mean) compensating variation is \$316 (\$477) for the efficiency decrease, and \$160 (\$250) for the electricity tax.

It is curious that these numbers are not identical. One explantation for the difference is households' ability to adjust the intensive margin in response to the tax, but not to the efficiency change. Buyers of new AC units will be worse off with a decrease in efficiency, and permanently so, since their marginal cost per BTU of cooling acts is now permanently higher. Further, the increase in the cost per BTU to AC purchasers due to the efficiency change is much more so than with a five percent electricity price increase. All AC users are affected by changes in electricity prices (creating a broad base), while only some are affected by efficiency changes (just much more intensely).

An analogous line of reasoning offers insights into the consumer welfare effects of an electricity tax versus efficiency standards. Efficiency standards will (at least in partial equilibrium) benefit consumers, while an electricity tax acts in the opposite direction. Costs of compliance to the efficiency standards would be incurred on the supply side (though the burden would be shared based

[37] Note that the decrease in efficiency of available AC units is transmitted to all households through their continuation value, since it reduces the utility garnered from potential future AC investment. This, however, is dwarfed by the utility impact on those who actually purchase a unit.

on the relative price elasticities of demand and supply). Since the model presented here includes only the simplest of supply sides (exogenous), a full exploration of these general equilibrium welfare effects is left to future work.

6 Conclusion

This paper examines how changes in attributes of energy-consuming durable goods affect the timing of purchase. Contrary to conventional wisdom, results indicate that consumers are forward-looking, and that expected future operating costs are a meaningful determinant of unit demand. The major difference between the model used in this study and that of counterparts in the literature is the role of consumer expectations. Here I explicitly allow consumers to be uncertain about the path of key product characteristics, whereas all other studies of demand for energy-consuming durables impose more restrictions on the extent of consumer rationality.

The importance of consumer expectations can be seen in two distinct ways after fitting the model to multiple years of the RECS micro-data. First, when testing the nature of consumer expectations explicitly, I reject myopia and (weakly) naïve expectations against the null hypothesis of rational expectations. Second, counterfactual simulations of changes in key product attributes reveal that consumer demand for air conditioners is more elastic with respect to energy efficiency than the up-front price of the durable.

How these results compare to Hausman's (1979) closely related work may be of particular interest to many readers. Despite his portrayal of consumers as more myopic, our findings are not necessarily inconsistent. Each approach sheds light on the relative importance of the trade-off between fixed and variable costs in household production technology, but our models speak to different dimensions. Hausman estimates the revealed tradeoff between up-front and ongoing costs *conditional on having decided to purchase*, a margin on which my approach is silent. On the other hand, my model sheds light on what product attributes induce a purchase to be made today (versus delaying), and what AC platform (room or central) is chosen. The apparent differences in the way consumers approach tradeoffs in the short run and long run highlight the need for a closer investigation. A model that combines a realistic framework of consumer expectations with data that includes more cross-sectional detail about durable good attributes would be a nice start, but is unfortunately not possible with the RECS data used in this paper.

Nonetheless, my findings have implications for both energy policy and future academic research on consumer demand for energy-consuming durable goods. Policies that seek to alter the timing of peoples' durable good decisions (e.g. automobile or appliance retirement programs, or weatherization subsidies) may be more cost-effective if they increase and emphasize the future stream of savings from increased energy efficiency, rather than focusing on the up-front cost. Further, when constructing empirical models of consumer behavior towards energy-consuming durables, researchers may want to move beyond the assumptions of perfect rental markets or myopia. If consumer expectations are important, these simplifications will miss important determinants of consumer decisions.

An odd feature of the results in this paper is the lack of equivalence between electricity price and energy efficiency effects on demand for air conditioners. A strictly neo-classical consumer ought to be indifferent between a permanent decrease in electricity prices and an equal increase in energy efficiency. The elasticities implied by the model here reveal consumers who are more responsive to a change in efficiency than a change in electricity price of equal magnitude. While this study provides no direct empirical tests of alternate explanations of this result, it is consistent with consumers viewing electricity price changes as transitory (whereas the model imposes permanence). An alternative explanation is more behavioral; at the time of purchase, electricity price is less salient than energy efficiency, and thus changes in the latter produce larger swings in demand. However, I leave these hypotheses to be tested in future research.

References

[1] Bento, A., L. Goulder, M. Jacobsen, and R. von Haefen (2009). "Distributional and Efficiency Impacts of Increased U.S. Gasoline Taxes." *American Economic Review*, Vol. 99, No. 3, 667-699.

[2] Brill, A., K. Hassett and G. Metcalf (1999). "Household Energy Conservation Investment and the Uninformed Consumer Hypothesis." Working Paper, Tufts University.

[3] Busse, M., C. Knittel, and F. Zettelmeyer (2009). "Pain at the Pump: How Gas Prices Affect Automobile Purchasing." Working paper.

[4] Dubin, J. and D. McFadden (1983). "A Heating and Cooling Load Model for Single-Family Detached Dwellings in Energy Survey Data." Working paper.

[5] Dubin, J. and D. McFadden (1984). "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica*, Vol. 52, No. 2, 345-362.

[6] Dubin, J., A. Miedema, and R. Chandran (1986). "Price Effects of Energy-Efficient Technologies: A Study of Residential Demand for Heating and Cooling." *The RAND Journal of Economics*, Vol. 17, No. 3, 310-325.

[7] Erdem, T., M.P. Keane, and B. Sun (2005). "Brand and Quality Choice Dynamics Under Price Uncertainty." *Quantitative Marketing and Economics*, Vol. 1, 5-64.

[8] Goldberg, P. (1995). "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry." *Econometrica*, Vol. 63, 891951.

[9] Gowrisankaran, G. and M. Rysman (2009). "Dynamics of Consumer Demand for New Durable Consumer Goods." Working paper.

[10] Hausman, J. (1979). "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *The Bell Journal of Economics*, Vol. 10, No. 1, 33-54.

[11] Hendel, I. and A. Nevo (2006). "Measuring the Implications of Sales and Consumer Stockpiling Behavior." *Econometrica*, Vol. 74, No. 6, 1637-1673.

[12] Ito, K. (2010). "How do Consumers Respond to Non-Linear Pricing? Evidence from Household Electricity Demand." Working paper.

[13] Knittel, C. and K. Metaxoglou (2011). "Estimation of Random Coefficient Demand Models: Challenges, Difficulties and Warnings." Working paper.

[14] Manning, W. (1998). "The Logged Dependent Variable, Heteroscedasticity, and the Retransformation Problem." *Journal of Health Economics*, Vol. 17, 283-295.

[15] Melnikov, O. (2001). "Demand for Differentiated Durable Products: The Case of the U.S. Computer Printer Industry." Working Paper.

[16] Prince, J. (2008). "Repeat Purchase amid Rapid Quality Improvement: Structural Estimation of the Demand for Personal Computers." *Journal of Economics and Management Strategy*, Vol. 17, No. 1, 1-33.

[17] Reiss, P. and M. White (2005). "Household Electricity Demand Revisited." *The Review of Economic Studies*, Vol. 72, 853-883.

[18] Rust, J. (1987). "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica*, Vol. 55, 999-1033.

[19] Schiraldi, P. (2011). "Automobile Replacement: A Dynamic Structural Approach." *The RAND Journal of Economics*, Vol. 42, No. 2, 266-291.

[20] Shin, J. (1985). "Perception of Price When Price Information Is Costly: Evidence from Residential Electricity Demand." *Review of Economics and Statistics*, Vol. 67, No. 4, 591-598.

[21] Tauchen, G. (1986). "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economic Letters*, Vol. 20, 177-181.

Tables and Figures

Table 1: Summary Statistics: RECS Sub-Sample (1990-2005)

	All US	US Non-CA	Northeast	Midwest	South	West
Income	48,595	47,860	57,341	45,324	44,184	52,996
	(36,081)	(35,643)	(39,127)	(32,200)	(34,904)	(38,285)
Square Footage	1,864	1,876	2,029	1,975	1,747	1,763
	(1,019)	(1,028)	(1,062)	(1,081)	(963)	(944)
Cooling Degree Days (65)	1,369	1,401	792	881	2,147	1,164
	(962)	(977)	(310)	(352)	(871)	(1,157)
Electricity Price ($/kWh)	0.09	0.09	0.11	0.09	0.07	0.10
	(0.12)	(0.13)	(0.08)	(0.22)	(0.04)	(0.06)
Kilowatt Hours Used	11,865	12,235	9,457	10,622	15,018	10,067
	(7,444)	(7,543)	(5,906)	(6,788)	(7,878)	(6,820)
Kilowatt Hours Used (AC Only)	1,815	1,921	773	1,225	3,348	794
	(2,251)	(2,299)	(1,107)	(1,293)	(2,679)	(1,625)
Electricity Expenditure	1,026	1,034	1,065	892	1,154	937
	(585)	(584)	(628)	(483)	(611)	(567)
Electricity Expenditure (AC Only)	152	159	106	110	241	92
	(189)	(192)	(152)	(123)	(221)	(172)
Fraction of HHs: No AC	0.267	0.247	0.373	0.187	0.076	0.575
Fraction of HHs: Room AC	0.267	0.280	0.390	0.262	0.277	0.124
Fraction of HHs: Central AC	0.466	0.473	0.237	0.551	0.647	0.301
N	11,261	10,491	2,423	2,971	3,669	2,198

Source: Energy Information Administration
Standard deviations in parentheses. All dollar values normalized to year 2000.

Table 2: Summary Statistics by Year: US non-CA

	1990	1993	1997	2001	2005
Income	50,978	41,317	46,069	52,062	45,920
	(39,623)	(32,371)	(34,945)	(37,637)	(28,335)
Square Footage	1,579	2,083	1,734	2,150	1,885
	(973)	(1,028)	(772)	(1,198)	(1,022)
Cooling Degree Days (65)	1,300	1,284	1,212	1,378	1,464
	(925)	(907)	(968)	(919)	(911)
Electricity Price ($/kWh)	0.06	0.06	0.08	0.10	0.13
	(0.08)	(0.02)	(0.03)	(0.03)	(0.04)
Fraction of HHs: No AC	0.313	0.284	0.219	0.201	0.162
Purchased Room AC (last 2yrs)	0.031	0.022	0.023	0.025	0.056
Fraction of HHs: Room AC	0.293	0.245	0.254	0.219	0.217
Purchase Central AC (last 2yrs)	0.045	0.051	0.054	0.064	0.070
Fraction of HHs: Central AC	0.319	0.398	0.450	0.491	0.496
N	2,653	2,464	2,149	1,754	1,471

Source: Energy Information Administration
Standard deviations in parentheses. All dollar values normalized to year 2000.

Table 3: Summary Statistics by AC Platform Type

	All US			US Non-CA		
	No AC	Room AC	Central AC	No AC	Room AC	Central AC
Income	42,986	41,620	54,882	40,612	41,290	54,432
	(34,356)	(33,384)	(37,106)	(32,069)	(33,242)	(37,054)
Square Footage	1,778	1,698	1,986	1,806	1,703	1,993
	(989)	(976)	(1,037)	(1,007)	(978)	(1,047)
Cooling Degree Days (65)	869	1,310	1,639	890	1,318	1,661
	(790)	(851)	(988)	(822)	(858)	(1,002)
Electricity Price ($/kWh)	0.10	0.09	0.09	0.09	0.09	0.08
	(0.08)	(0.23)	(0.04)	(0.09)	(0.23)	(0.04)
Kilowatt Hours Used	8,334	10,479	14,252	8,760	10,575	14,564
	(5,512)	(6,592)	(7,795)	(5,775)	(6,639)	(7,844)
Kilowatt Hours Used (AC Only)	0	1,206	2,990	0	1,228	3,092
	0	(1,421)	(2,422)	0	(1,435)	(2,450)
Electricity Expenditure	834	946	1,158	851	949	1,156
	(484)	(543)	(617)	(488)	(546)	(612)
Electricity Expenditure (AC Only)	35	110	230	42	111	233
	(110)	(137)	(205)	(120)	(138)	(208)
N	5,247	3,002	3,012	4,959	2,592	2,940

Source: Energy Information Administration

Standard deviations in parentheses. All dollar values normalized to year 2000.

Figure 1: Room AC Price and Energy Efficiency

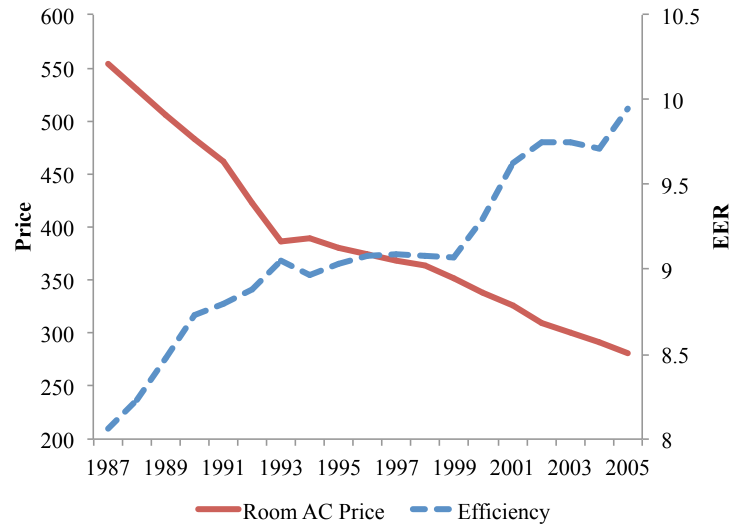

Figure 2: Central AC Price and Energy Efficiency

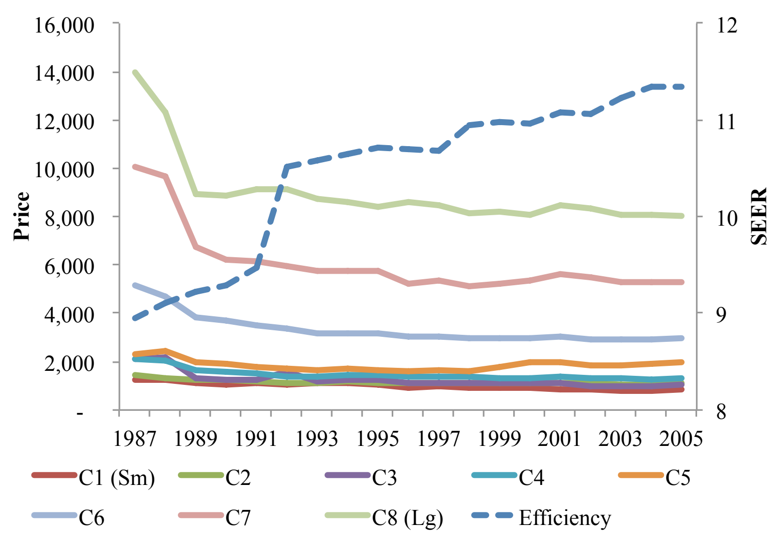

Notes: Prices normalized to year 2000 dollars. Energy efficiency is sales-weighted. SEER equals the ratio of cooling output in BTU to the power consumption in Watts/hr, evaluated over an entire season. EER is a cooling rate equal to BTUh/Watts, evaluated at an outdoor temperature of 95 degrees.

Figure 3: Evolution of AC Stock: 1990-2005

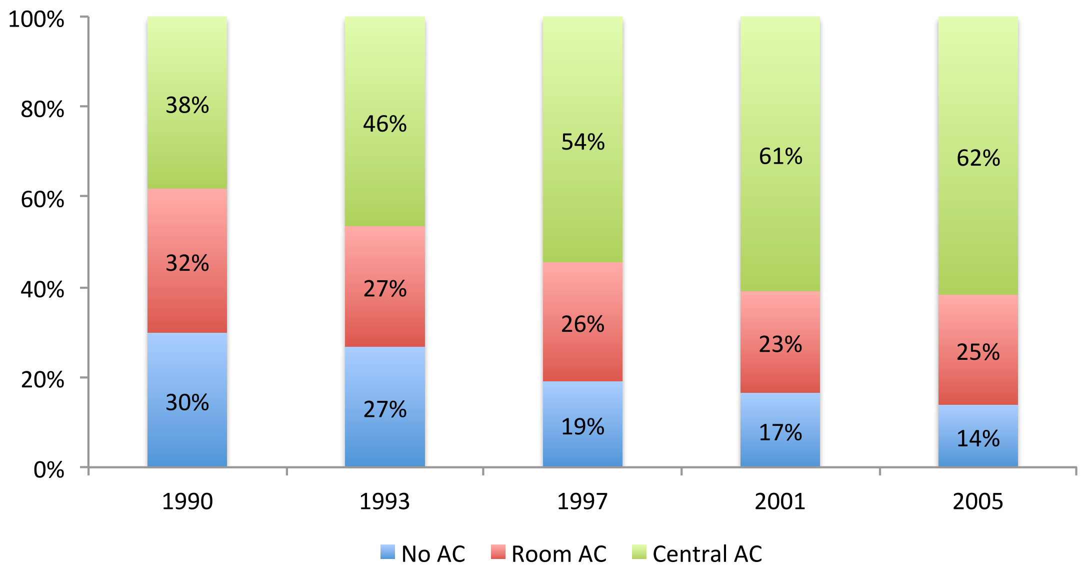

Note: This statistics are derived from the weighted subsample of each repeated cross-section, as described in the Data section.

Table 4: Marginal Effects: Central AC (kwh)

	All US		US Non-CA	
Electricity Price	-0.170	***	-0.068	*
	(0.039)		(0.038)	
Cooling Degree Days	0.989	***	0.961	***
	(0.031)		(0.031)	
Square Footage	0.376	***	0.334	***
	(0.041)		(0.033)	
Income	0.206	***	0.207	***
	(0.024)		(0.025)	
N	5,247		4,959	

*Standard errors in parentheses. Statistical significance at p<0.1, p<0.05, and p<0.01 denoted by *, **, and *** respectively. Energy efficiency coefficient fixed at -0.925 according to estimate of the rebound effect. All variables in logs.*

Table 5: Marginal Effects: Room AC (kwh)

	All US		US Non-CA	
Electricity Price	-0.339	***	-0.337	***
	(0.050)		(0.056)	
Cooling Degree Days	1.074	***	1.092	***
	(0.051)		(0.052)	
Square Footage	-0.020		-0.023	
	(0.057)		(0.055)	
Income	0.114	***	0.126	***
	(0.036)		(0.036)	
2nd Room Unit Indicator	0.530	***	0.520	***
	(0.063)		(0.065)	
3rd Room Unit Indicator	0.191	**	0.183	**
	(0.087)		(0.092)	
N	3,012		2,940	

*Standard errors in parentheses. Statistical significance at p<0.1, p<0.05, and p<0.01 denoted by *, **, and *** respectively. Energy efficiency coefficient fixed at -0.925 according to estimate of the rebound effect. All variables in logs.*

Table 6: Structural Model Parameter Estimates

	US Non-CA	
Comfort	5.934	***
	(1.684)	
AC Expenditure	-0.336	***
	(0.008)	
Room Platform Dummy	0.959	***
	(0.064)	
Central Platform Dummy	1.118	†
N	10,491	

Bootstrapped standard errors in parentheses.
† Calibrated to match aggregate transition probabilities.
*Statistical significance at p<0.1, p<0.05, and p<0.01 denoted by *, **, and *** respectively.*

Table 7: Cumulative Elasticities

Cumulative Elasticity wrt Efficiency

	2007	2009	2011	2013	2015
Central Unit Demand	0.975 ***	0.919 ***	0.861 ***	0.806 ***	0.755 ***
	(0.267)	(0.200)	(0.189)	(0.219)	(0.206)
Room Unit Demand	0.273 **	0.248 **	0.247 **	0.236 **	0.228 **
	(0.114)	(0.109)	(0.102)	(0.095)	(0.090)
Electricity Usage	-0.119 **	-0.172 ***	-0.219 ***	-0.263 ***	-0.303 ***
	(0.056)	(0.055)	(0.055)	(0.057)	(0.058)

Cumulative Elasticity wrt AC Unit Price

	2007	2009	2011	2013	2015
Central Unit Demand	-0.245 ***	-0.243 ***	-0.241 ***	-0.245 ***	-0.248 ***
	(0.047)	(0.039)	(0.034)	(0.031)	(0.028)
Room Unit Demand	-0.118 *	-0.116 *	-0.124 **	-0.126 **	-0.127 **
	(0.062)	(0.061)	(0.060)	(0.059)	(0.058)
Electricity Usage	0.004 ***	0.005 ***	0.007 ***	0.008 ***	0.010 ***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)

Cumulative Elasticity wrt Electricity Price

	2007	2009	2011	2013	2015
Central Unit Demand	-0.024	-0.024	-0.024	-0.024	-0.024
	(0.150)	(0.115)	(0.094)	(0.083)	(0.076)
Room Unit Demand	-0.347 *	-0.349 *	-0.325 *	-0.265	-0.220
	(0.208)	(0.197)	(0.185)	(0.173)	(0.165)
Electricity Usage	-0.697 ***	-0.709 ***	-0.716 ***	-0.721 ***	-0.725 ***
	(0.191)	(0.191)	(0.193)	(0.194)	(0.194)

*Bootstrapped standard errors in parentheses. Statistical significance at $p<0.1$, $p<0.05$, and $p<0.01$ denoted by *, **, and *** respectively.*

Table 8: Period-by-Period Elasticities

Period-by-period Elasticity wrt Efficiency

	2007	2009	2011	2013	2015
Central Unit Demand	0.975 ***	0.863 ***	0.747 ***	0.644	0.552 ***
	(0.267)	(0.152)	(0.242)	(0.535)	(0.143)
Room Unit Demand	0.273 **	0.218 **	0.243 **	0.175	0.159
	(0.114)	(0.109)	(0.110)	(0.134)	(0.138)
Electricity Usage	-0.119 **	-0.224 ***	-0.314 ***	-0.393 ***	-0.463 ***
	(0.056)	(0.055)	(0.055)	(0.062)	(0.063)

Period-by-period Elasticity wrt AC Unit Price

	2007	2009	2011	2013	2015
Central Unit Demand	-0.245 ***	-0.240 ***	-0.237 ***	-0.259 ***	-0.256 ***
	(0.047)	(0.030)	(0.022)	(0.019)	(0.019)
Room Unit Demand	-0.118 *	-0.114 *	-0.148 ***	-0.141 **	-0.131 *
	(0.062)	(0.061)	(0.057)	(0.065)	(0.070)
Electricity Usage	0.004 ***	0.007 ***	0.010 ***	0.013 ***	0.016 ***
	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)

Period-by-period Elasticity wrt Electricity Price

	2007	2009	2011	2013	2015
Central Unit Demand	-0.024	-0.024	-0.024	-0.024	-0.024
	(0.150)	(0.083)	(0.077)	(0.064)	(0.056)
Room Unit Demand	-0.347 *	-0.351 *	-0.248	0.075	0.137
	(0.208)	(0.211)	(0.249)	(0.375)	(0.433)
Electricity Usage	-0.697 ***	-0.722 ***	-0.730 ***	-0.735 ***	-0.739 ***
	(0.191)	(0.197)	(0.198)	(0.197)	(0.196)

*Bootstrapped standard errors in parentheses. Statistical significance at $p<0.1$, $p<0.05$, and $p<0.01$ denoted by *, **, and *** respectively.*

Table 9: Robustness Tests: Implied Elasticities

Robustness Scenario:	Base Case	1	2	3	4	5
β:	0.9	0.9	0.9	0.9	0.9	0.95
γ^c/γ^r:	0.7	0.7	0.7	0.5	1	0.7
Rebound:	0.075	0	0.15	0.075	0.075	0.075
Efficiency						
Central AC purchases	0.975	1.015 *	1.435 *	1.136 *	1.073 *	1.170 *
Room AC purchases	0.273	0.174 *	-0.159	0.086 *	0.199 *	0.350 *
Electricity usage	-0.119	-0.129 *	-0.187 *	-0.161 *	-0.114 *	-0.139 *
Unit Price						
Central AC purchases	-0.245	-0.250 *	-0.238 *	-0.236 *	-0.259 *	-0.212 *
Room AC purchases	-0.118	-0.102 *	-0.017 *	-0.044 *	-0.119 *	-0.122 *
Electricity usage	0.004	0.004 *	0.005 *	0.005	0.003 *	0.004 *
Electricity Price						
Central AC purchases	-0.024	-0.017 *	-0.070 *	-0.034 *	0.011 *	-0.047 *
Room AC purchases	-0.347	-0.419 *	-0.313 *	-0.388 *	-0.535 *	-0.483 *
Electricity usage	-0.697	-0.694 *	-0.501 *	-0.384 *	-0.365 *	-0.699 *

*Note: * denotes an elasticity estimate that falls within the 95 percent confidence interval of the Base Case Estimate*

Table 10: Matching Aggregate Central AC Moment – Replacements Fraction of All Purchases

	US Non-CA
Actual	0.53
Predicted	0.55
	(0.02)

Bootstrapped standard deviations in parentheses

Table 11: Unconditional Choice Probabilities – Observed and Predicted

	No Purchase		Buy Room AC		Buy Central AC	
Year	Observed	Predicted	Observed	Predicted	Observed	Predicted
1990	0.922	0.883	0.031	0.060	0.046	0.057
1993	0.926	0.895	0.020	0.041	0.054	0.064
1997	0.921	0.889	0.023	0.032	0.055	0.079
2001	0.904	0.897	0.025	0.026	0.071	0.077
2005	0.878	0.896	0.047	0.025	0.075	0.079

Note: Observed and predicted probabilities are calculated using RECS sampling weights. Predicted values are generated using the optimal structural coefficient estimates.

Table 12: Model Fit

Likelihood Ratios		
H_0: Rational Expectations		
	H_1: Myopic	H_1: Naïve
US	6,865.8 ***	15.0 *
US Non-CA	4,981.8 ***	6.0

*Statistical significance at p<0.1, p<0.05, and p<0.01 denoted by *, **, and *** respectively.*

Table 13: Compensating Variation

	5.0% Electricity Price Inrease	31.6% Decrease in Market Energy Efficiency*
Median	\$159.6	\$315.9
Mean	\$250.2	\$476.7
	(227.9)	(471.8)

**Aggregate consumer welfare-equivalent to a 5% electricity price increase.*
Standard deviations in parentheses

A Aggregate State Transitions

I calculate micro moments in the RECS data that I then use to match the predicted aggregate state transition probabilities implied by the likelihood estimates. First I determine the fraction of central AC purchases that are replacements. All central AC purchases are either replacements or upgrades from room AC. Comparing the fraction of the population that purchases central in a given period to the change in the fraction of the population that has central at the beginning and end of that period distinguishes between the two types of purchases. Let λ_t^{+c} represent fraction of the total population purchasing a central AC unit in period t, s_t^c represent the fraction of all households ending period t with a central unit installed, and μ_t be the fraction of households that purchased a central unit as a replacement (i.e. that started period t in the central AC state). I calculate μ_t as follows:

$$\mu_t = 1 - \frac{s_t^c - s_{t-1}^c}{\lambda_t^{+c}}.$$

Having calculated the fraction of central replacement purchases, I next estimate probabilities associated with originating states of households that purchased central as an upgrade. Under the transition assumptions in the model, each of these households began the period with either one, two, or three room units installed. A similar non-parametric approach as used for μ_t could identify the probability of upgrading to central from each of the room platform states; however, one might also imagine that certain household characteristics (e.g. square footage) might be correlated with the number of room units. In order to incorporate heterogeneity and allow the probabilities of originating states to differ across households, I estimate an ordered probit model. Specifically, I limit my sample for this exercise to households with room air conditioners, and use household characteristics as variables to explain the number of room units chosen. I then use the parameter estimates to assign predicted probabilities of each originating room AC state to all households that purchased a central unit.