Journal of Environmental Economics and Management 68 (2014) 141–160

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

journal homepage: www.elsevier.com/locate/jeem

Durable goods and long-run electricity demand: Evidence from air conditioner purchase behavior

David Rapson

Department of Economics, University of California, Davis, One Shields Ave, Davis, CA 95616, United States

ARTICLE INFO

Article history:
Received 4 December 2012
Available online 5 May 2014

JEL classification:
C61
L68
Q40
Q54

Keywords:
Dynamic discrete choice
Durable goods
Energy demand

ABSTRACT

I estimate a dynamic structural model of demand for air conditioners, the most energy-intensive home appliance in the US. The model explores the links between demand for durable goods and expected changes in key attributes: energy efficiency and price. I incorporate expectations explicitly as a feature of the choice setting, and use parameter estimates from the model to calculate durable good demand elasticities with respect to energy efficiency, electricity price, and price of the durable itself. These estimates fill a large gap in the literature, and also shed light on consumer behavior in this setting. Results indicate that consumers are forward-looking and value the stream of future savings derived from energy efficiency.

Introduction

Dynamic considerations are important in a variety of decision contexts relating to the environment. Depletion of common-pool resources, extraction of non-renewable resources, and stock externalities are just a few of the most vivid examples where actions taken today alter tomorrow's state variables and choice set, with environmental consequences.[1] An extensive literature explores economic efficiency and optimal policy in these settings, yet in only rare instances are dynamic elements of decisions explicitly built into empirical models of consumer choice. In this paper I estimate a dynamic discrete choice model of the demand for air conditioners, the most energy-intensive home durable good in the United States and a major contributor to residential sector carbon emissions. The model explicitly allows for different behavioral hypotheses with respect to the nature of consumer expectations. The underlying rationale for using such a model is that consumers do not simply choose *whether or not* to purchase an energy-efficient durable, but also *when* to do so.[2] A main result of this study is that, after accounting for this feature of consumer behavior, the data fit more closely with a model of forward-looking rationality than with myopia or naïve expectations.[3]

E-mail address: dsrapson@ucdavis.edu

[1] For classic references, see Gordon (1954) on common pool resources, Hotelling (1931) on natural resource extraction, and Keeler et al. (1972) on optimal pollution control. Each spawned a long literature, recent contributions to which include Smith et al. (2009) on fisheries, Lin (2013) on dynamics of oil extraction, and Nordhaus (1991) and Falk and Mendelsohn (1993) on optimal abatement paths.

[2] Jaffe and Stavins (1995) emphasize this distinction in a model of energy-efficient technology diffusion, and it is a key feature of a large body of dynamic discrete models used in Industrial Organization, beginning with Rust (1987).

[3] Under "myopia", consumers fully discount future utility, whereas naïve expectations imply that the current value of key attributes is also the expected future value.

http://dx.doi.org/10.1016/j.jeem.2014.04.003
0095-0696/

The purpose of this paper is in part to advocate for the inclusion of dynamic discrete choice methods in the empirical toolkit for applied environmental economists. In addition, the application itself is relevant today given the urgent focus on choosing between environmental policy instruments to combat global climate change. 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).[4] 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. It is thus reasonable to seek to quantify the extent to which these standards achieve their intended effect. As a first step to understanding whether efficiency standards reduce energy usage (a necessary but not sufficient condition for economic efficiency), one must understand how consumer demand responds to changes in energy efficiency, the relevant product attribute.

In the model, fully informed and rational consumers form expectations about future changes in key product attributes (e.g. energy efficiency and price), and time their purchases in order to maximize the present value of the associated stream of net benefits. During the period studied here (1987–2005), 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 and 49 percent less expensive. The dynamic framework encapsulates intensity-of-use (and the future stream of production costs that it implies) that varies with energy efficiency and energy prices. The model builds on dynamic discrete choice literature, drawing primarily on Rust (1987).[5] Rust estimates the optimal replacement timing of bus engines, a decision that has several similarities to the timing of consumer durable good purchases. I use a repeated cross section of appliance purchase decisions which classifies air conditioners into vintages. The model is then extended to incorporate a wide gamut of observed consumer heterogeneity. Individual level data on demand makes the empirical approach similar to Goldberg (1995), extended to include dynamic considerations.

I use structural parameter estimates from the model to calculate a variety of policy-relevant summary statistics, including the long-run elasticities of demand for air conditioners with respect to durable price, energy efficiency, and the price of electricity. In this and other environmental settings, estimates of these statistics help us to alleviate a scarcity in the empirical literature. The structural parameters are also then used to simulate the model forward under different assumptions about state variable transitions. In other words, I simulate policy counterfactuals which allow for predictions of environmental policy effectiveness. The counterfactuals reflect a Pigouvian tax, increases in average energy efficiency, and a durable good subsidy.[6]

I also use the model to generate suggestive evidence about the nature of inter-temporal consumer preferences. The discount rate is non-parametrically unidentified (Magnac and Thesmar, 2002), making it impossible to reliably estimate directly. However, one may compare model fit under various specifications as evidence of forward-looking behavior, which I do by comparing information criteria (AIC and BIC). I estimate the model under three behavioral hypotheses relating to consumer expectations over the paths of dynamic state variables (energy efficiency and AC unit prices). The *rational expectations* hypothesis assumes that consumers generate expectations over these variables that are ex post correct. At the opposite end of the dynamic behavior spectrum is the *myopic* consumer, who exhibits an infinite discount rate. Finally, the intermediate case modeled here is that of *naïve expectations*, under which consumers expect dynamic state variables to follow a random walk. These are admittedly extreme versions of the behavioral possibilities.

This study has five main qualitative findings. First, the model provides evidence that behavior governed by rational expectations fits the data better than behavior under the alternatives (myopic and naïve expectations). The difference in model fit is equivalent to roughly seven additional degrees of freedom. Second, consumers generally exhibit more elastic demand with respect to energy efficiency than with respect to either the durable price or the price of electricity. This is true for central AC units without exception, whereas efficiency-elasticity for room AC units is roughly the same as the electricity price elasticity in the short run. While economic theory dictates that consumers should respond symmetrically to changes in energy efficiency and electricity prices, I allow these effects to differ in the model. This is consistent with other energy demand studies[7] and is motivated by the observation that consumers actually appear to respond differently to economic incentives on these margins. Third, increases in the average energy efficiency of air conditioners raise air conditioner demand and lower 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. Energy efficiency gains outweigh energy increases from first-time purchases 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.

Poor model fit under the myopia hypothesis combines with high estimates of the demand elasticity with respect to energy efficiency to suggest that consumers are forward-looking in the timing of their choice of durable. This deviates from

[4] There may also be information asymmetries leading to principal–agent problems, as estimated in Gillingham et al. (2012).

[5] A similar approach has been used in the environmental literature to examine fishing location choice (Provencher and Bishop, 1997; Smith, 2005).

[6] One might consider this study to be a demand-side analog to Jaffe and Stavins (1995), who examine dynamic supply-side response to these same policy alternatives. They make the reasonable assertion that the long-run development and adoption of energy efficient technologies is inextricably linked to the extent of environmental protection induced by policy.

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

the conventional view of consumers as myopic, although evidence of forward-looking behavior continues to mount.[8] 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.

This study extends and contributes to the literature in two main ways. First, it extends models of demand for energy-consuming durable goods to reflect dynamic aspects of choice. Dubin and McFadden (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. Each of these analyses is thus better-suited to reflect short-run trade-offs. 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.

The second contribution is a richer understanding of the nature of consumer choice, and in particular the availability of long-run demand elasticities which are important inputs to several vibrant strands of the environmental economics and policy literature. There are startlingly few empirical estimates of the elasticity of demand for durable goods, including home appliances, and those that exist use decades-old data. Dale and Fujita (2008) review these studies, and add their own estimates for refrigerators, clothes washers and dishwashers. For room air conditioning, Golder and Tellis (1998) apply a diffusion model to data on shipments from 1946 to 1962, and estimate the own price elasticity of demand to be -0.37.

My estimates of own-price elasticity for air conditioners are smaller – approximately -0.12 for room units and -0.24 for central air conditioners – and may be due to a number of potential factors, including different modeling choices (e.g. dynamics). The main point, though, is that the literature is bereft of estimates of these important statistics. Integrated assessment models, such as those used by the EPA and other institutions, require these elasticities as inputs, as do efforts to forecast climate change adaptation scenarios along the lines of Auffhammer and Aroonruengsawat (2011) or Mansur et al. (2008). Long-run forecasting exercises such as these will be aided by a more accurate understanding of consumer behavior (e.g. long-run demand elasticities), and there is much room still to add to this effort.

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 1990–2005.[9] 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.

The units of measure for $E_{it} \in \{E_{it}^r, E_{it}^c\}$ 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 kW h) is converted into cooling energy (British Thermal Units or BTUs). 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.[10] Tables 1–3 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. California is consistently an outlier on energy policy. To avoid concerns that its inclusion would

[8] Busse et al. (2013) show that consumers respond to high gas prices by purchasing more fuel-efficient cars.

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

[10] 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.

Table 1
Summary statistics: RECS sub-sample (1990–2005).

Variable	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	1864	1876	2029	1975	1747	1763
	(1019)	(1028)	(1062)	(1081)	(963)	(944)
Cooling degree days (65)	1369	1401	792	881	2147	1164
	(962)	(977)	(310)	(352)	(871)	(1157)
Electricity price ($/kW h)	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	9457	10,622	15,018	10,067
	(7444)	(7543)	(5906)	(6788)	(7878)	(6820)
Kilowatt hours used (AC only)	1815	1921	773	1225	3348	794
	(2251)	(2299)	(1107)	(1293)	(2679)	(1625)
Electricity expenditure	1026	1034	1065	892	1154	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.

Variable	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	1579	2083	1734	2150	1885
	(973)	(1028)	(772)	(1198)	(1022)
Cooling degree days (65)	1300	1284	1212	1378	1464
	(925)	(907)	(968)	(919)	(911)
Electricity price ($/kW h)	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 2 yrs)	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 2 yrs)	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	2653	2464	2149	1754	1471

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

contaminate the analysis with systematic unobserved heterogeneity, I eliminate it from my preferred subsample. For this reason, I display statistics for this segment in the "US Non-CA" column of Table 1.

In my analysis, there are three platform types: central, room, and no AC. "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 split systems have separate indoor evaporator and outdoor condenser units. Households with central AC tend to have a higher income, have a larger home, live in a warmer climate, and have a lower electricity price than do 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).[11] 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. 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 of electricity demand and the

[11] 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, $high_t$, in degrees fahrenheit and, in the case of the data used here, 65.

Table 3
Summary statistics by AC platform type.

Variable	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	1778	1698	1986	1806	1703	1993
	(989)	(976)	(1037)	(1007)	(978)	(1047)
Cooling degree days (65)	869	1310	1639	890	1318	1661
	(790)	(851)	(988)	(822)	(858)	(1002)
Electricity price ($/kW h)	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	8334	10,479	14,252	8760	10,575	14,564
	(5512)	(6592)	(7795)	(5775)	(6639)	(7844)
Kilowatt hours used (AC only)	0	1206	2990	0	1228	3092
	0	(1421)	(2422)	0	(1435)	(2450)
Electricity expenditure	834	946	1158	851	949	1156
	(484)	(543)	(617)	(488)	(546)	(612)
Electricity expenditure (AC only)	35	110	230	42	111	233
	(110)	(137)	(205)	(120)	(138)	(208)
N	5247	3002	3012	4959	2592	2940

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

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.[12]

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. If consumers know where the kink points are in their price schedule, the quantity of 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 that extent consumers respond to price, they do so to average price, not marginal. This is the assumption made, for example, by Dubin 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 (2014) lends empirical support for the use of average prices in this setting.

Variation in the data that identifies purchase timing 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, one must model the household's AC characteristics before the purchase (the “originating state”). This procedure is discussed in the Model section.

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.[13] I discuss the role of transition paths in detail in the Model section.

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. As such, I require data on annual average energy efficiency, price, and quantity sold.

[12] The state variables in my model are also likely included in the EIA estimation procedure. This is of course not ideal, but is unavoidable.
[13] See the Appendix for details.

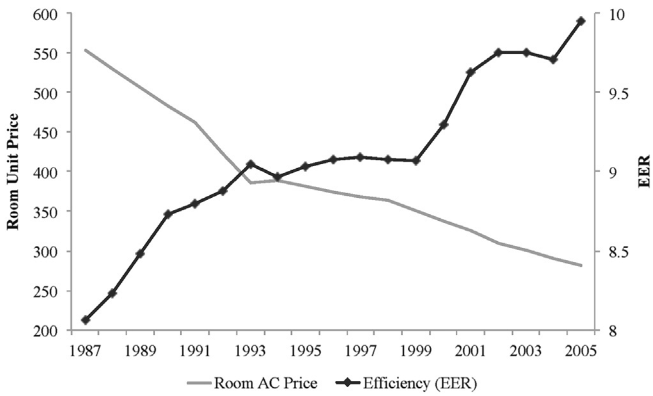

Fig. 1. Room AC price and energy efficiency. *Notes*: Prices normalized to year 2000 dollars. Energy efficiency is sales-weighted. SEER equals the ratio of cooling output in BTU to the power consumption in W/h, evaluated over an entire season. EER is a cooling rate equal to BTU h/W, evaluated at an outdoor temperature of 95 degrees.

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. Fig. 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 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. Fig. 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.

A noteworthy aspect of Fig. 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 and helps us 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.[16]

Fig. 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 explores this relationship.

Model

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 from cooling services. 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

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

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

[16] 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 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 moving 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).

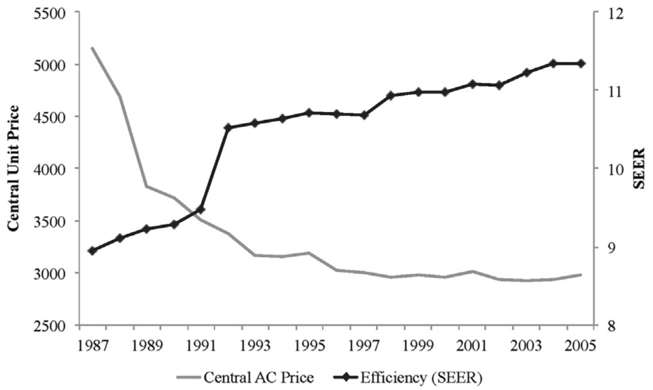

Fig. 2. Central AC price and energy efficiency.

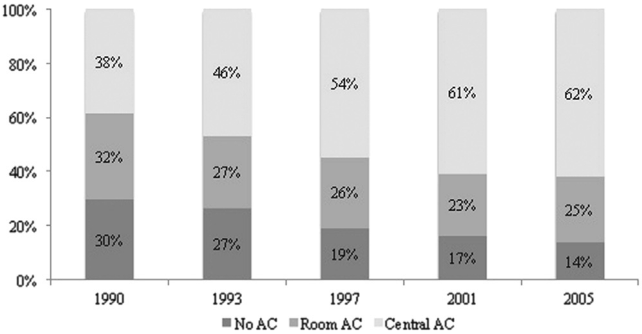

Fig. 3. Evolution of AC stock: 1990–2005. *Note*: This statistics are derived from the weighted subsample of each repeated cross-section, as described in the Data section.

a purchase, its stock of the durable is augmented according to the characteristics of the good. Since characteristics of the durables for sale in the market change over time, households may elect to defer their purchase until technology has progressed. Each period, they choose between doing nothing (i.e., wait to purchase in a future period), buying a room AC unit, or installing/replacing the central AC unit.

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 ($Z_{it} \in \{0, r_1, r_2, r_3, c\}$). These states are mutually exclusive.[17] The productive potential of durables persists across periods, but depreciates deterministically. There are no secondary markets, and no transaction costs of purchasing in the 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 of units available for purchase 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).

[17] In the data, some households have both room and central air conditioners concurrently. However, only 2.3 percent 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 0–4, 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.

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 $A(Z_{it}) \in \{0, 1, 2\}$ differ 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, reflect the “hotness” of climate relative to 65 degrees.[18] 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.

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 the use of room units unless every room in the structure is occupied.[19] 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, Eqs. (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,

$$k_{it} = \underset{k}{\operatorname{argmax}}\, u_{it}\ (k_{it}|a_{it}), \tag{4}$$

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

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

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

conditioning platform

$$k_{it}^r = b_0^r + b_E^r E_{it}^r + b_\Phi^r \Phi_{it}^r + \epsilon_{it}^r \quad (5)$$

$$k_{it}^c = b_0^c + b_E^c E_{it}^c + b_\Phi^c \Phi_{it}^c + \epsilon_{it}^c \quad (6)$$

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, T_{it}.

Durable good purchase decision

I assume that utility is log-linear in cooling utility, δ_{it}, and AC-related expenditures, $p_{it}^e k_{it}$ and P_t^j. This yields the following functional forms:

$$u_{it}(\hat{k}_{it}|a_{it}) = \left\{ \begin{array}{ll} \delta_{it}(\widehat{k_{it}}) + \alpha_1 \ln(p_{it}^e \widehat{k_{it}}) + \varepsilon_{it} & \text{if} a_{it} = 0 \\ \delta_{it}(\widehat{k_{it}}) + \alpha_1 \ln(p_{it}^e \widehat{k_{it}} + P_t^r) + \varepsilon_{it} & \text{if } a_{it} = 1 \\ \delta_{it}(\widehat{k_{it}}) + \alpha_1 \ln(p_{it}^e \widehat{k_{it}} + P_t^c) + \varepsilon_{it} & \text{if } a_{it} = 2 \end{array} \right\} \quad (7)$$

and

$$\delta_{it} = \alpha_2 \ln T_{it}(\widehat{k_{it}}|a_{it}) + \alpha_3 D_{it}^r + \alpha_4 D_{it}^c \quad (8)$$

Upgrading one's air conditioning holdings changes derived electricity demand differently for households with room versus central units. 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.

Eqs. (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. The utility from cooling of those with no AC (the "outside option" in this model) is equal to $(-\alpha_2 \ln C_i)$, reflecting the fact that 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 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.

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 in the data 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 can be done using the results of the household production optimization above (which predicts optimal cooling and electricity usage conditional on household attributes).

Another 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. Given the data exclude new construction, it is unlikely that a household desiring cooling services would jump straight to a central AC retrofit of the home without first availing itself of a simple, inexpensive, and timely option of purchasing a room unit. The model also constrains households to purchase only one AC unit in a single period. The maximum number of room units allowed in the model is three.[20] Households that have already adopted central AC have the choice of replacing their current unit with a new one or doing nothing. This restriction is common in the discrete choice literature, and makes the model tractable while still allowing for a rich set of stock evolution patterns.

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 $\rho = 0.01$.[21] 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. 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. This deterioration implies that older units use more electricity to generate a given amount of cooling than new units. This holds even if there was 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.

[20] A mere 2.3 percent of households in the data have more than three room units.

[21] 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.

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

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

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) and others, the solution to this value function requires integration over the realizations of two random variables: $\varepsilon_{it}(=\varepsilon(Z_{it}))$ and Ω_{it}. I follow Rust (1987) in assuming "conditional independence" between these two variables, which is satisfied implicitly by assuming that ε_{it} is *iid* Extreme Value Type I. One can then separately consider the expectation of the value function conditional on 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}.[22]

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

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:

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

Since all $a \in A(Z_{it})$ are mutually exclusive and exhaustive, the probability of a household choosing not to buy is

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

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 (1987), 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. 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 ($4 \times 4 \times 3 \times 2$ for CDD, electricity price, square footage, and income, respectively). The dimension of each household's dynamic state space depends on its endowment of air conditioner platform. Households with no AC have 9 potential price and efficiency states that correspond to a purchase of room AC (3×3 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 the 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)), and the value function associated with each of these stylized households must be calculated.

To generate a continuous likelihood function, and to map the data into model space, I implement a piecewise linear spline. That is, I determine the discrete bins that are closest to each observation in the data, and use linear interpolation to determine the observation-specific value function under each element of the choice set. I then construct the probability of each choice as in (12). This procedure 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.

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

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

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.

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.[23]

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

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

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), $\hat{e}_t$, transition probabilities can be calculated for the range spanned by each bin of the dynamic variable, as in Tauchen (1986). The resulting transition matrices operationalize the expectation in Eq. (11) in discretized state space.

Identification

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 Figs. 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 Fig. 2.

The second source of identifying variation 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.

Identifying derived demand for cooling

Identification requirements for demand on the intensive margin are standard. Consistent estimation of Eqs. (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.[24] 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

[23] Some studies (originally Melnikov, 2001, and later Gowrisankaran and Rysman, 2009; Schiraldi, 2011, and Hendel and Nevo, 2006) reduce multiple dynamic variables into a (logit) inclusive value, which they then assume follows a first-order Markov process.

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

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.

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.

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:

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

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.

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 energy efficiency improvements 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.

First-stage derived demand analysis

The derived demand for cooling electricity is estimated using Eqs. (5) and (6). Cross-sectional variation in price, efficiency, and household attributes pins 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

Table 4
Marginal effects (kW h).

Variable	Central AC		Room AC	
	All US	US non-CA	All US	US non-CA
Electricity price	−0.165***	−0.056*	−0.342***	−0.339***
	(0.024)	(0.034)	(0.014)	(0.039)
Cooling degree days	0.999***	0.967***	1.079***	1.094***
	0.040	0.017	0.070	0.012
Square footage	0.378***	0.330***	−0.019	−0.022
	0.034	0.032	0.055	0.087
Income	0.206***	0.207***	0.115***	0.127***
	0.033	0.014	0.024	0.033
WW2 Dum			0.525***	0.515***
			0.021	0.081
WW3 Dum			0.207***	0.199
			0.062	0.130
N	5247	4959	3012	2940

Standard errors in parentheses. Energy efficiency coefficient fixed at −0.925 according to estimate of the rebound effect. All variables in logs.
* Statistical significance at $p < 0.1$.
*** Statistical significance at $p < 0.01$.

estimates from the literature. Under the assumption that this restriction is correct, and that consumers respond to average (not marginal) electricity prices, Eqs. (5) and (6) will yield consistent estimates of derived demand for cooling electricity.

Marginal effects corresponding to the electricity demand analysis (Eqs. (5) and (6)) are presented in Table 4. Geographically, the 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.[25] 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 Eqs. (5) and (6). Table 4 presents the marginal effects of each variable (in logs) on derived electricity demand from cooling. 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 estimates from the literature (e.g. Reiss and White, 2005; Ito, 2014; Jessoe and Rapson, 2014). 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).

Structural model results and implications

Estimates of the model's structural parameters (α_1–α_4 from Eqs. (7) and (8)) are displayed in Table 5. 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 likelihood estimator for each draw.

[25] 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.

Table 5
Structural model parameter estimates.

Variable	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[a]
N	10,491

Bootstrapped standard errors in parentheses.
[a] Calibrated to match aggregate transition probabilities.
*** Statistical significance at $p<0.01$.

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.[26] 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 are, however, very good, with a mean predicted fraction of replacements of 0.55 (with standard deviation 0.02), relative to 0.53 in the data.

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 compare the goodness-of-fit of the rational expectations hypothesis against the two alternative and less-sophisticated hypotheses of consumer behavior (myopia and naïve expectations) by examining standard statistics. The benefits of this exercise are two-fold: it aids us in interpreting the structural model and counterfactual results, and provides further justification for the dynamic structural approach to demand estimation in this setting.

Table 9 shows the annual unconditional choice probabilities in the data and those predicted by the model. From 1991 to 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.[27] 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 compare model fit under the assumption of rational expectations with the two alternate behavioral hypotheses: myopic consumers and naïve expectations. This is done by imposing each alternative behavioral hypotheses on the model, solving for the optimal parameters, then examining statistical measures of goodness-of-fit (log-likelihood, AIC, and BIC). Recall that "rational expectations" describes consumers who form expectations over the path of variables that are ex post correct. "Myopic" consumers 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.

Table 10 displays statistics that emerge from each of these models. Higher values of the log-likelihood and smaller values of AIC and BIC indicate superior model fit. The rank order of goodness-of-fit favors rational over naïve expectations, and naïve expectations over myopia.[28] When comparing the log-likelihood under rational and naïve expectations, the difference (8 for the entire US and 37 for U.S. non-CA) indicates a substantial improvement in explanatory power. A χ-square statistic comparing these log-likelihoods is large enough to support seven additional degrees of freedom under rational expectations. The myopia hypothesis generates an even larger difference in model fit (poorer), which ought not be surprising given the nature of the choice setting. These results suggest both that the dynamic features of the setting are important and that consumers exhibit forward-looking tendencies. However, there is admittedly a gradient of potential

[26] The process by which the aggregate moments are calculated is described in the Appendix.
[27] 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.
[28] The rank order is identical across these measures, since the number of parameters and sample size are constant across behavioral specifications.

Table 6
Cumulative elasticities.

Elasticity Type	2007	2009	2011	2013	2015
Cumulative elasticity wrt efficiency					
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					
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					
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 $p < 0.1$.
** Statistical significance at $p < 0.05$.
*** Statistical significance at $p < 0.01$.

Table 7
Period-by-period elasticities.

Elasticity Type	2007	2009	2011	2013	2015
Period-by-period elasticity wrt efficiency					
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					
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					
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 $p < 0.1$.
** Statistical significance at $p < 0.05$.
*** Statistical significance at $p < 0.01$.

forward-looking sophistication, and the three hypotheses modeled here represent some of the more extreme possibilities. Given the challenges of estimating intertemporal preferences (e.g. Magnac and Thesmar, 2002), there remains much scope for future research on the topic.

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 in Tables 6 and 7, which are discussed in the following subsection. 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 Eq. (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 subsequent periods. 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.

To evaluate the effects of price changes and energy efficiency improvements, 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.[29]

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 sampling draws as 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).[30] 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. All elasticities should be interpreted as market elasticities, as opposed to elasticities for individual products.[31]

Table 6 displays cumulative elasticities of demand for air conditioners and electricity with respect to energy efficiency, unit price, and electricity price. These statistics are calculated at two-year intervals when simulating demand forward ten years. 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 7. I focus the discussion that follows on the cumulative results, due to their greater policy-relevance.

The high elasticities of demand for AC units 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 energy efficiency improvements (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

[29] 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

$$\varepsilon_T^{kE} = \frac{\frac{dk_T}{k_T}}{\frac{dE}{E}}.$$

The cumulative analog is

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

[30] See Erdem et al. (2005) and Hendel and Nevo (2006).

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

Table 8
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[a]	1.435[a]	1.136[a]	1.073[a]	1.170[a]
Room AC purchases	0.273	0.174[a]	−0.159	0.086[a]	0.199[a]	0.350[a]
Electricity usage	−0.119	−0.129[a]	−0.187[a]	−0.161[a]	−0.114[a]	−0.139[a]
Unit price						
Central AC purchases	−0.245	−0.250[a]	−0.238[a]	−0.236[a]	−0.259[a]	−0.212[a]
Room AC purchases	−0.118	−0.102[a]	−0.017[a]	−0.044[a]	−0.119[a]	−0.122[a]
Electricity usage	0.004	0.004[a]	0.005[a]	0.005	0.003[a]	0.004[a]
Electricity price						
Central AC purchases	−0.024	−0.017[a]	−0.070[a]	−0.034[a]	0.011[a]	−0.047[a]
Room AC purchases	−0.347	−0.419[a]	−0.313[a]	−0.388[a]	−0.535[a]	−0.483[a]
Electricity usage	−0.697	−0.694[a]	−0.501[a]	−0.384[a]	−0.365[a]	−0.699[a]

[a] Denotes an elasticity estimate that falls within the 95 percent confidence interval of the Base Case Estimate.

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 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 rebates associated with the EnergyStar program), 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 unit price incentives 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 effects that act in different directions. All households will move along their demand curve and consume less on the intensive margin. On the extensive margin, demand for new, efficient units should increase as inefficient units are replaced. 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. Results from the present model indicate that the latter effect dominates, and that unit demand decreases in electricity prices (though only slightly, and statistically indistinguishable from zero for central AC). This leads to larger (in absolute value) price elasticity of demand for electricity, as seen in the bottom row in Tables 7 and 8. While the (short run) estimates of the price elasticity of demand for electricity in the literature reside in the range −0.1 to −0.5 (Reiss and White, 2005; Ito, 2014; Jessoe and Rapson, 2014), the long-run estimates, particularly in Table 7, are closer to −0.7 to −0.8. Estimates of the price elasticity derived from static frameworks neglect the effect of electricity price changes on the extensive margin, and a strength of the current approach to capture the joint effect.

These elasticity estimates are robust to changes in assumptions about some of the key input parameters. Table 8 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.

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

Table 9
Unconditional choice probabilities – observed and predicted.

Year	No purchase		Buy room AC		Buy central AC	
	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 10
Model fit.

Behavioral hypothesis	US			US non-CA		
	LogL	AIC	BIC	LogL	AIC	BIC
Rational expectations	−3717	7426	7471	−3517	7027	7072
Naïve expectations	−3725	7442	7487	−3554	7100	7145
Myopia	−7150	14,293	14,338	−6990	13,973	14,018
k	4	4	4	4	4	4
N	11,261	11,261	11,261	10,491	10,491	10,491

Note: Models differ by the hypothesized nature of consumer expectations (and are thus non-nested), but are parameterized identically.

(consumer-borne) electricity tax to a change 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.[32]

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 11 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 energy efficiency decrease and \$160 (\$250) for the electricity tax.

It is noteworthy 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 energy efficiency standards. *Ceteris paribus*, efficiency standards will 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 on the relative price elasticities of demand and supply). Since the model presented here includes only the simplest supply side (exogenous), a full exploration of these general equilibrium welfare effects is left to future work.

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

[32] 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.

Table 11
Compensating variation.

Statistic	5.0 percent electricity price increase	31.6 percent decrease in market energy efficiency[a]
Median	\$159.6	\$315.9
Mean	\$250.2	\$476.7
	(227.9)	(471.8)

[a] Aggregate consumer welfare-equivalent to a 5 percent electricity price increase. Standard deviations in parentheses.

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 examining the nature of consumer expectations, the data fit a rational expectations model more closely than myopia or naïve 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.[33] Nonetheless, this paper contributes in two meaningful ways to the literature. The modeling exercise demonstrates how dynamic consumer behavior can be considered when calculating demand elasticities of durable goods; and the elasticities themselves are also important as inputs into integrated assessment models and other efforts to simulate the long run effects of climate change and adaptation.

An interesting 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.

Acknowledgments

I thank Marc Rysman, Max Auffhammer, Scott Carrell, Randall Ellis, Ken Gillingham, Matthew Harding, Jordi Jaumandreu, Chris Knittel, Laurence Kotlikoff, Michael Manove, Robert Margo, J. Bradford Rice, Jim Sanchirico, Jeremy Smith, and Robert Stavins for their comments. Excellent research assistance during earlier revisions was provided by Tom Blake. This work has also benefited 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.

Appendix 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).

[33] 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 data used in this paper.

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 one, two, or three room units installed. A similar non-parametric approach to that 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.

References

Auffhammer, M., Aroonruengsawat, A., 2011. Simulating the impacts of climate change, prices and population on Californias residential electricity consumption. Clim. Change 109, S191–S210.

Bento, A., Goulder, L., Jacobsen, M., von Haefen, R., 2009. Distributional and efficiency impacts of increased U.S. gasoline taxes. Am. Econ. Rev. 99 (3), 667–699.

Brill, A., Hassett, K., Metcalf, G., 1999. Household Energy Conservation Investment and the Uninformed Consumer Hypothesis. Working Paper. Tufts University.

Busse, M., Knittel, C., Zettelmeyer, F., 2013. Are consumers myopic? Evidence from new and used car purchases. Am. Econ. Rev. 103 (1), 220–256.

Dale, L., Fujita, S., 2008. An Analysis of the Price Elasticity of Demand for Household Appliances. Working Paper. Lawrence Berkeley National Laboratory.

Dubin, J., McFadden, D., 1984. An econometric analysis of residential electric appliance holdings and consumption. Econometrica 52 (2), 345–362.

Dubin, J., Miedema, A., Chandran, R., 1986. Price effects of energy-efficient technologies: a study of residential demand for heating and cooling. RAND J. Econ. 17 (3), 310–325.

Erdem, T., Keane, M.P., Sun, B., 2005. Brand and quality choice dynamics under price uncertainty. Quant. Mark. Econ. 1, 5–64.

Falk, I., Mendelsohn, R., 1993. The economics of controlling stock pollutants: an efficient strategy for greenhouse gases. J. Environ. Econ. Manag. 25, 76–88.

Gillingham, K., Harding, M., Rapson, D., 2012. Split incentives in residential energy consumption. Energy J. 33 (2), 37–62.

Golder, P., Tellis, G., 1998. Beyond diffusion: an affordability model of the growth of new consumer durables. J. Forecast. 17, 259–280.

Goldberg, P., 1995. Product differentiation and oligopoly in international markets: the case of the U.S. automobile industry. Econometrica 63, 891–951.

Gordon, H.S., 1954. The economic theory of a common-property resource: the fishery. J. Polit. Econ. 62 (2), 124–142.

Gowrisankaran, G., Rysman, M., 2009. Dynamics of Consumer Demand for New Durable Consumer Goods. Working Paper.

Hausman, J., 1979. Individual discount rates and the purchase and utilization of energy-using durables. Bell J. Econ. 10 (1), 33–54.

Hendel, I., Nevo, A., 2006. Measuring the implications of sales and consumer stockpiling behavior. Econometrica 74 (6), 1637–1673.

Hotelling, H., 1931. The economics of exhaustible resource extraction. J. Polit. Econ. 29 (2), 137–175.

Ito, K., 2014. How do Consumers Respond to Non-Linear Pricing? Evidence from Household Electricity Demand. Am. Econ. Rev. 104 (2), 537–563.

Jaffe, A., Stavins, R., 1995. Dynamic incentives of environmental regulations: the effects of alternate policy instruments on technology diffusion. J. Environ. Econ. Manag. 29, S43–S63.

Jessoe, K., Rapson, D., 2014. Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use. Am. Econ. Rev. 104 (4), 1417–1438.

Keeler, E., Spence, M., Zeckhauser, R., 1972. The optimal control of pollution. J. Econ. Theory 4, 19–34.

Lin, C., 2013. Strategic decision-making with information and extraction externalities: A structural model of the multi-stage investment timing game in offshore petroleum production. Rev. Econ. Stat. 95 (5), 1601–1621.

Magnac, T., Thesmar, D., 2002. Identifying Dynamic Discrete Decision Processes. Econometrica 70 (2), 801–816.

Manning, W., 1998. The logged dependent variable, heteroscedasticity, and the retransformation problem. J. Health Econ. 17, 283–295.

Mansur, E., Mendelsohn, R., Morrison, W., 2008. Climate change adaptation: a study of fuel choice and consumption in the US energy sector. J. Environ. Econ. Manag. 55 (2), 175–193.

Melnikov, O., 2001. Demand for Differentiated Durable Products: The Case of the U.S. Computer Printer Industry. Working Paper.

Nordhaus, W., 1991. To slow or not to slow: the economics of the greenhouse effect. Econ. J. 101 (407), 920–937.

Provencher, B., Bishop, R., 1997. An estimable dynamic model of recreation behavior with an application to great lakes angling. J. Environ. Econ. Manag. 33, 107–127.

Reiss, P., White, M., 2005. Household electricity demand revisited. Rev. Econ. Stud. 72, 853–883.

Rust, J., 1987. Optimal replacement of GMC bus engines: an empirical model of Harold Zurcher. Econometrica 55, 999–1033.

Schiraldi, P., 2011. Automobile replacement: a dynamic structural approach. RAND J. Econ. 42 (2), 266–291.

Shin, J., 1985. Perception of price when price information is costly: evidence from residential electricity demand. Rev. Econ. Stat. 67 (4), 591–598.

Smith, M., 2005. State dependence and heterogeneity in fishing location choice. J. Environ. Econ. Manag. 50, 319–340.

Smith, M., Sanchirico, J., Wilen, J., 2009. The economics of spatial-dynamic processes: applications to renewable resources. J. Environ. Econ. Manag. 57, 104–121.

Tauchen, G., 1986. Finite state Markov-chain approximations to univariate and vector autoregressions. Econ. Lett. 20, 177–181.