INSTITUTIONS AND INCENTIVES FOR THE EFFICIENT USE OF ENERGY

By

Dylan Brewer

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Economics – Doctor of Philosophy
Environmental Science and Policy – Dual Major

2019

ABSTRACT

INSTITUTIONS AND INCENTIVES FOR THE EFFICIENT USE OF ENERGY

By

Dylan Brewer

Chapter 1: Moral Hazard and Equilibrium Sorting in Residential Energy Markets

This chapter studies tenant behavior in rental housing when the landlord pays for heating. I develop a model in which renters have heterogeneous preferences for home size and indoor temperature. When energy is costly, renters choose smaller apartments and turn down the heat—or sort into apartments with landlord-pay energy bills. I estimate the model using exogenous variation in energy prices and use a machine-learning algorithm to explore preference heterogeneity. Surprisingly, I find that renters who prefer hotter temperatures do not systematically choose landlord-pay units, though I am unable to rule out sorting on preferences for unobserved home attributes. Eliminating moral hazard by forcing all renters to pay their own bill reduces energy consumption by 25% due to renters turning down the heat (22%) and choosing smaller units (3%). Moral hazard in residential energy contracts cost the United States \$839 million per year in welfare losses including \$238 million from carbon emissions.

Chapter 2: The Remarkably Inelastic Demand for Home Heating Services: A Choice Experiment

I conduct a stated-choice experiment that presents research subjects with hypothetical costs to adjust their thermostats. I estimate responses to the cost of heating and analyze the causes for heterogeneity in household demand for energy services, using the experimental results as a complete-information baseline. I find that even at the highest price level, half of the participants exhibit zero response to price. Further, I find that participants' experimental behavior with complete information can explain the full range of observed real-world temperature settings, suggesting a limited role for informational barriers or price salience issues in energy-service demand heterogeneity. Individuals with higher temperature preferences are more price responsive, suggesting that increasing block

pricing for energy may reduce energy consumption while minimizing the regressivity of the energy pricing program.

Chapter 3: Negawatts v. Megawatts: Demand Response in Wholesale Electricity Markets (with Katherine Wagner)

We study the Federal Energy Regulatory Commission's (FERC) wholesale demand-response program. In this program, participants sell a reduction in electricity use from a baseline on the wholesale electricity market as if it was electricity generation. We show that demand-response participants have different incentives to consume electricity based on whether they participate directly, through a third-party aggregator, or through a utility aggregator. We argue that the current wholesale demand-response program is inefficient and should be replaced by retail demand-response programs run by utilities.

ACKNOWLEDGEMENTS

I thank Soren Anderson for his careful mentorship, friendship, and genuine interest in my well-being. I would not be the same economist without Soren's advice and guidance. In addition, I thank Joseph Herriges, Kyoo il Kim, and Scott Swinton for serving on my committee, reading my work, and passing on their wisdom. I would also like to thank Steven Haider for his professional mentorship, Cody Orr for many brilliant insights, and Michael Conlin for getting me to go the extra mile. Special thanks to Lee Coppock for introducing me to economics and inspiring this journey.

My time spent at the Property and Environment Research Center in Bozeman, Montana was critical to the development of this thesis. I would like to thank the many individuals that provided feedback on my research while I was there: Randal Rucker, V. Kerry Smith, Timothy Fitzgerald, Maximilian Auffhammer, Matthew Kahn, Roger Meiners, and Ryan Kellogg. I also want to thank the office staff at PERC, in particular Sharie Rucker and Dianna Reinhart, for their excellent support.

The Environmental Science and Policy Program at Michigan State University provided funding for the choice experiment in the thesis. I would like to thank Joseph Hamm for his help in the experimental design. I am grateful for fellowship support from the Department of Economics at Michigan State University, the Graduate School at Michigan State University, the Environmental Science and Policy Program at Michigan State University, the Institute for Humane Studies, and the Property and Environment Research Center while earning my PhD. I appreciate the office staff of the Department of Economics, and in particular Lori Jean Nichols, for their excellent support of my endeavors while at Michigan State.

Thank you to Katherine Wagner for her important work as co-author of Chapter 3. I am grateful for comments by seminar participants at Michigan State University, the CU Environmental & Resource Economics Workshop, Midwest Energy Fest (Northwestern), Energy Camp at Haas, Camp Resources (NC State), the Midwest Economics Association Annual Meeting, the Heartland Workshop (University of Illinois-Urbana Champaign), Davidson College, James Madison University, Wake Forest University, the University of Georgia, and Georgia Tech.

Finally, thank you to my family for their support, advice, and love. I especially want to thank my future wife Amanda for being my best friend, supporting my dreams, reading my drafts, becoming an honorary economist, and moving with me to Georgia. We did it!

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

MORAL HAZARD AND EQUILIBRIUM SORTING IN RESIDENTIAL ENERGY MARKETS

1.1 Introduction

Economic wisdom suggests that individuals make socially efficient decisions only when they face the full costs and benefits of those decisions. In landlord-tenant rental regimes, the energy-efficiency and energy-use decisions are fundamentally divided in what is called the "split-incentive problem." When the landlord pays for utilities, tenants face no marginal cost for energy use and so are expected to use more energy than if they were facing the full marginal cost—a form of moral hazard.[1] Meanwhile, household sorting exacerbates incentive problems through the choice of payment regime and home attributes that affect energy efficiency. First, a household that expects to consume more energy services will have a higher willingness to pay for landlord-pay utilities, all else equal. Second, conditional on choosing a landlord-pay unit a household may choose a larger home because it knows it will not pay the full cost to heat the larger home.[2] Thus, the relevant causal effect of landlord-pay utilities on energy use includes both the direct effect via temperature setting as well as the indirect effect via the choice of payment regime and housing attributes. An estimated 17 percent of US rental housing has landlord-pay heating contracts, suggesting that eliminating moral hazard may provide large private energy savings and reduced external costs from fossil fuel combustion (US Census, 2013).

In this paper, I develop and estimate a model of renters making joint home-efficiency and energy-use decisions that characterize moral hazard and equilibrium sorting in the residential energy market.

[1]Conversely, when the tenant pays for utilities, landlords do not directly benefit from energy efficiency and so are expected to under-invest in energy efficiency—also a form of moral hazard.

[2]Recent work on energy use suggests that individuals are cognizant of energy costs and savings. Myers (2017) provides convincing evidence that when purchasing a home, households discount future energy costs at 8-10 percent. In commercial real estate, properties built under more stringent energy-efficiency standards rent and sell at a premium relative to those built under less stringent energy-efficiency standards (Papineau, 2017).

I find that landlord-pay households spend 25 percent more on energy for heating than they would if all rental contracts required the household to pay for energy.[3] The moral hazard effect dominates, contributing to 22 percent of extra heating expenditures per household. Meanwhile, landlord-pay households choose housing units that are 140 square feet larger, contributing an additional 3 percent of extra heating expenditures. In addition, I find that for a fixed cost of heating, households that choose tenant-pay contracts prefer higher thermostat settings than households that choose landlord-pay contracts as predicted by observables. This counter-intuitive result suggests that selection into the landlord-pay regime is driven by preferences for unobserved characteristics of landlord-pay units. Finally, I provide evidence of energy-efficiency economies of scale in multi-unit apartment buildings, suggesting a possible explanation for the persistence of landlord-pay utility regimes.

In the model, renters have heterogeneous preferences for home size and temperature setting. When energy is costly, renters either choose smaller apartments and turn down the heat to conserve energy, or they sort into apartments with landlord-pay energy bills. Using data on housing rents, temperature settings, and exogenous variation in energy prices, I estimate equilibrium hedonic prices for home size and landlord-pay heating, the engineering energy cost of home size, and the engineering energy cost of winter temperature settings. Given these estimates, I then estimate how an average household's indoor temperature setting responds to the estimated marginal cost of heating the home. I corroborate the results using a choice experiment that I designed and implemented on a nationally representative sample of US households. These parameter estimates allow me to construct unique household utility functions from the model that I use to simulate counterfactuals. I simulate the effects of eliminating moral hazard by requiring that all tenants pay for their own energy bills and allow tenants to respond by re-sorting into housing units of different sizes. With this counterfactual, I quantify the energy impacts of moral hazard, selection into housing regime, and selection of housing attributes in the landlord-pay utility regime.

This paper makes several contributions to the economics literature studying energy efficiency. This is the first paper to use a structural model to analyze the split-incentive problem for residential

[3] This calculation is relative to the average bill for tenant-pay units because the actual heating bill is not observed for landlord-pay households.

energy contracts. Experimental interventions requiring tenants to pay the bills find average energy savings as high as 20 and 25 percent (Dewees and Tombe, 2011; Elinder et al., 2017), while quasi-experimental approaches measure energy savings of less than 3 percent (Levinson and Niemann, 2004; Jessoe et al., 2018).[4] I measure an average impact consistent with the literature, but my structural approach provides new insights into the relative sizes of moral hazard and equilibrium sorting. This paper complements theoretical and empirical work studying underinvestment in energy efficiency due to split-incentive problems (e.g. Jaffe and Stavins, 1994; Harding et al., 2000; Davis, 2012; Gillingham et al., 2012). While recent reviews conclude that the "energy-efficiency gap" is small in most sectors (Allcott and Greenstone, 2012; Gillingham and Palmer, 2014), recent work on energy use in rental housing suggests that asymmetric information between landlords and tenants significantly distorts energy-efficiency investment incentives (Myers, 2018). Further, I add to recent work investigating energy-efficiency scale economies in urban space (Borck and Brueckner, 2017). The modeling and estimation strategy used here can be adapted to other contexts with an up-front energy-efficiency choice followed by a downstream energy-use choice.

In addition, this paper connects the literature on equilibrium sorting in housing markets (Kuminoff et al., 2013) to the literature estimating moral hazard and adverse selection under asymmetric information.[5] I show how to estimate two-stage partial-equilibrium models of selection and moral

[4]Estimates of conservation behavior by tenants who pay for utilities vary by context. Levinson and Niemann (2004) estimate that paying for heat reduces energy expenditures by less than 1 percent on average. On the other extreme, Elinder et al. (2017) find 25 percent average reductions in electricity use when Swedish renters were experimentally switched to tenant-pay. The top 20 percent of energy users completely drive this result; the majority of households in the study did not change behavior. Furthermore, these households did not use electricity for heating, and heat use is highly inelastic. An experiment based on a single condominium complex found that submetering decreased electricity consumption by 20 percent (Dewees and Tombe, 2011). Jessoe et al. (2018) find that switching to tenant-pay results in a 2.9 percent reduction in average daily electricity use among the top decile of commercial users, but that the effect for the remaining 90 percent of users is negligible. The difference between single-site experimental and multi-site quasi-experimental treatment effects may reflect site-selection bias on the experimental side (Allcott, 2015).

[5]Most recent papers in this literature focus on asymmetric information in health insurance markets; e.g., Brot-Goldberg et al. (2017); Finkelstein et al. (2016); Baicker et al. (2015); Autor et al. (2014); Einav et al. (2013). Other applications study lending markets (e.g. Veiga and Weyl, 2016; Crawford et al., 2018), crop insurance markets (He et al., 2017a,b), online marketplaces (Hui

hazard using an assumption on the parametric form of the utility function as in Bajari and Benkard (2005). One key insight from this paper is that the first-stage choice of home attributes affects the second-stage heating costs. Households' choices of housing and subsequent heating use reveal relative preferences for housing attributes and energy services when housing and fuel marginal costs vary over time and space. The main appeal of this approach is that it allows for the estimation of completely heterogeneous utility functions without making assumptions on the distribution of the error term and with minimal supply-side assumptions. Another benefit is that estimation is computationally simple and transparent. The parametric assumption on the utility function provides closed-form solutions for the unknown parameters of each household's utility function. These parameters can be evaluated using household consumption decisions, hedonic prices, and downstream energy costs estimated from the data.

Finally, I introduce a novel machine-learning approach to estimate energy-preference parameter heterogeneity in the structural model. Households that pay their own energy bills do not reveal the counterfactual winter temperature settings they would choose if the landlord was to pay for heat. Households in landlord-pay units do reveal their "bliss point" temperature preferences (i.e., the temperature chosen at zero marginal cost). This bliss point temperature setting is likely a function of physiological characteristics such as age and sex.[6] I train a machine-learning algorithm to predict underlying bliss-point temperature settings using demographic covariates and the revealed bliss points of households in landlord-pay units. These predictions characterize household heating demand when there is zero marginal cost for temperature setting. I use these predictions to explore heterogeneity in household energy demand, contributing to work analyzing household energy demand heterogeneity (e.g. Reiss and White, 2005; Auffhammer and Rubin, 2018).

et al., 2016), worker contracts (Jackson and Schneider, 2011, 2015), and vehicle leasing (Weisburd et al., 2018).

[6] Appendix B.1 discusses the physiological determinants of temperature setting preferences with a thorough review of science and engineering studies.

1.2 Data and descriptive evidence of moral hazard

In this section, I provide descriptive evidence of moral hazard in rental energy use. In 2013, an estimated 17 percent of US renters had landlord-pay heating (US Census, 2013).[7] Space heating accounts for approximately 42 percent of residential energy use by American households and is the single largest non-transportation residential energy item (EIA, 2013). For this reason, I focus on space heating and winter thermostat setting to analyze heterogeneous demand for energy services and moral hazard in residential energy contracts.

Moral hazard arises in the landlord-pay regime because households do not face the marginal cost of energy use. A housing unit's utility payment regime has a larger effect on the household's marginal heating cost than home size or any other housing attribute. If utility payments are bundled into the rent, the household has an effective marginal cost of *zero* for heating services regardless of how large the home is. When the marginal cost of heating is zero, the household will heat their home to a "bliss point," regardless of energy price or outside temperature. Thus, for a given household, heating use will be weakly higher in a landlord-pay regime than in a tenant-pay regime. Landlord-pay households do not respond to market signals of scarcity or to the increased cost of heating due to low outside temperatures.

Sorting into landlord-pay regimes on temperature preferences may also occur. Given heterogeneous preferences for heat, households will differ in willingness to pay for the zero-marginal-cost heating offered in a landlord-pay regime. Households with higher bliss point temperature will be willing to pay more to heat without marginal cost all else equal. Those with high energy demand are thus likely to sort into landlord-pay units. This sorting results in higher energy use than would occur under randomly assigned regimes.

To test for moral hazard, I turn to data in the Residential Energy Consumption Survey (RECS). The RECS asks landlord-pay and tenant-pay renters about their winter temperature settings. The

[7] This paper focuses on the rental case. While some owner-occupied housing units have the utility payment bundled into the homeowners' association fee, this case is relatively rare. Two percent of homeowners had utility payments bundled into the homeowners' association fee.

RECS is a repeated cross-sectional survey of nationally representative US households collected every four years. I examine 3,993 electric-heating and gas-heating renters from the 2001, 2005, and 2009 surveys—3,164 tenant-pay households and 829 landlord-pay households.[8] In addition to thermostat settings, the survey collects heating bills for tenant-pay units, housing characteristics, heating and cooling degree days, location at the census division level, and household demographics.

Figure A.1 shows that landlord-pay households set their thermostats higher during the winter than tenant-pay households during the winter. This difference by regime in average thermostat setting can be explained either by moral hazard or by selection into landlord-pay regime on temperature preference. Table A.1 shows differences in the demographics of tenant-pay and landlord-pay households. Landlord-pay renters are older and poorer than tenant-pay renters—and thus may also be expected to differ in temperature preferences. However, figure A.2 shows a larger difference in temperature setting between the two regimes when the household is out of the home. This evidence suggests that some of the difference in temperature setting between regimes is driven exclusively by moral hazard rather than sorting.

Households also endogenously select housing attributes that influence energy cost. For example, perhaps an additional 100 square feet increases the rent by five dollars each month but also increases energy cost by a dollar each month. Then the true cost of 100 additional square feet is six dollars per month rather than five. Even if households take energy prices as given, a household knows that it can influence the cost of energy services by selecting different housing attributes. Likewise, it is more expensive to increase the thermostat setting in a larger housing unit than a smaller unit.

I investigate the relationship between housing attributes, rent price, and utility payment regime using data from the American Housing Survey (AHS). The AHS is a panel of housing units labeled at the Metropolitan Statistical Area (MSA) level and includes information on rental price and unit characteristics for 106,062 rental units from 1997 to 2013.[9] Table A.2 displays sample means for landlord-pay and tenant-pay units. On average, landlord-pay units are cheaper, smaller, older,

[8] See the appendix for notes on how the final sample is chosen.

[9] An MSA is a statistical region often centered around a city and the surrounding population centers.

and in larger buildings than tenant-pay units. In addition, landlord-pay units are less likely than tenant-pay units to have large energy-using appliances such as in-unit laundry and air conditioning. Thus, landlord-pay units are older but otherwise have attributes associated with lower energy use. It is difficult to draw conclusions about the impacts of equilibrium sorting in this context, because if more landlord-pay units are in high-cost rental areas we may be mistaking a movement along the demand curve of an attribute for a systematic difference in regime type. The role of home attribute prices becomes clear in the formal model—where I turn to next.

1.3 Modeling and empirical strategy

In this section, I develop a model that describes how renters make joint choices of housing and heating. Renters with heterogeneous preferences select housing units with differentiated attributes and utility payment regimes that influence the cost of heating. Renters subsequently choose temperature setting, trading off heating with other consumption. I begin with a general model that demonstrates the connection between housing attributes, heating cost, and temperature setting. Next, I specify parametric forms of the utility function, heating cost function, and hedonic price function that allow me to isolate and estimate the welfare impacts of moral hazard and equilibrium sorting.

1.3.1 General model

Consider households $i \in \{1, ..., I\}$ choosing to live in housing units $j \in \{1, ..., J\}$ located in a housing market. Housing units differ in basic housing characteristics x_j (e.g., square feet), and the utility payment regime R_j (a binary variable equal to one if the tenant pays the marginal cost of utilities and zero if the landlord pays). The equilibrium per-month rental price of each housing unit is a function of attributes and is given by $p_j = p(x_j, R_j, \xi_j)$, where x_j are observed housing characteristics and ξ_j are researcher-unobserved housing characteristics. In this model, the choice of which market to live in is made before considering housing and energy use so that a rental unit in one housing market is not a substitute for a rental unit in another housing market. For example, households choose jobs that determine which housing market is available to them before considering housing choice.[10]

Given incomes y_i, households derive utility from the housing characteristics x_j and ξ_j, consumption of an outside good z_i, and the temperature setting within the housing unit s_i. I represent household preferences with the utility function $u_i(x_j, s_i, \xi_j, z_i)$, which is concave in x_j, s_i, and ξ_j. Furthermore, there exists a bliss point temperature preference s_i^b such that

$$\left.\frac{\partial u_i}{\partial s}\right|_{s=s_i^b} = 0, \text{ and} \tag{1.1}$$

$$\left.\frac{\partial^2 u_i}{\partial s^2}\right|_{s=s_i^b} < 0. \tag{1.2}$$

Intuitively, the bliss point s_i^b is an individual's preferred heating temperature when facing a zero marginal cost for heating. In addition, I adopt the assumptions of Bajari and Benkard (2005) to guarantee the existence of a differentiable hedonic price surface in equilibrium.[11]

The amount the household spends on heating is a function of temperature setting s_i, housing characteristics x_j, the ambient temperature in the heating season T_j, and the market price of energy P_e. Let $H(x_j, s_i, T_j, P_e, \xi_j)$ be the expected per-month cost of heating unit j to temperature setting s given outdoor temperature T_j and energy prices P_e. $H(\cdot)$ can be understood as a heating cost function. Households sign leases without exact knowledge of the future outside temperature. Households choose among units j with different housing attributes x_j and ξ_j, taking into consideration future heating costs. Once the households have moved in, they observe the outside temperature and choose the temperature setting s_i. At this point, the households cannot change their optimally chosen housing units $j^*(i)$, so the housing attributes are fixed and denoted x_{j^*}.

[10] This assumption is realistic in modeling the short run where a household's employment is fixed and therefore tied to a housing market. In the long run, a household may move to take advantage of differential housing and energy prices across markets. The short run view is more realistic given large moving costs.

[11] Specifically, the assumptions on the utility function are that (1) u_i is continuously differentiable in z_i and strictly increasing in z_i with $\partial u_i / \partial z_i > \epsilon$ for some $\epsilon > 0$ and $z_i > 0$, (2) u_i is Lipschitz continuous in the housing attributes, and (3) u_i is strictly increasing in ξ_j. Given discrete housing attributes, condition (2) is not satisfied, but preferences over continuous attributes are identified as the choice set becomes large (Bajari and Benkard, 2005).

In this framework, each household maximizes utility by first choosing a housing unit and then choosing a temperature setting. Each household solves this problem by reverse induction: each decides how much heat to use in a housing unit with given characteristics x_j and ξ_j and housing regime R_j and then uses this planned behavior to choose the preferred housing unit $j^*(i)$. The subsequent heating choice is a budget-constrained trade-off between consumption of energy services and consumption of an outside good z_i with price equal to one.

$$\begin{aligned} \text{Heating choice: } \max_{\{s_i, z_i\}} \quad & u_i(x_{j^*}, s_i, \xi_{j^*}, z_i) \\ \text{s.t. } \ & y_i \geq p(x_{j^*}, R_{j^*}, \xi_{j^*}) + R_{j^*} \cdot H(x_{j^*}, s_i, T_{j^*}, P_e, \xi_{j^*}) + z_i. \end{aligned} \tag{1.3}$$

Note that the household cannot change the housing characteristics x_{j^*} because it has already chosen its housing unit $j^*(i)$. This problem yields the first-order condition for maximization:

$$\implies [s_i] : \quad \frac{\partial u_i}{\partial s_i} \bigg/ \frac{\partial u_i}{\partial z_i} = R_{j^*} \cdot \frac{\partial H}{\partial s_i}. \tag{1.4}$$

Equation 1.4 describes the household's temperature setting behavior. The left side of the equation is the marginal rate of substitution between energy services and consumption of the outside good. The right side is the price ratio with the price of the outside good normalized to one. If the household is in a landlord-pay regime ($R_{j^*} = 0$), the household sets temperature to the bliss point s_i^b regardless of cost. Intuitively, the bliss point is the temperature at which the utility function peaks—beyond this point, a person feels too hot and before this point a person feels too cold. Denote $s_i = s_i^*(x_{j^*}, R_{j^*}, T_{j^*}, P_e, y_i)$ as the optimal "home heating rule" as determined by the first-order condition in equation 1.4.

When sorting into a housing unit j, households consider the rent price of the unit and future heating costs:

$$\begin{aligned} \text{Housing choice: } \max_{j} \quad & u_i(x_j, s_i^*, \xi_j, z_i) \\ \text{s.t. } \quad & y_i \geq p(x_j, R_j, \xi_j) + R_j \cdot H(x_j, s_i, T_j, P_e, \xi_j) + z_i. \end{aligned} \tag{1.5}$$

I treat future energy prices as known by households, which is a realistic assumption given regulated residential energy prices. I further assume that utility is linearly separable in the outside good to allow me to interpret the opportunity cost of housing and heating in dollar terms.[12] If there is a continuum of available housing attributes such that there is continuity in attributes x and ξ, the first-order conditions to this problem characterize sorting behavior:

$$\implies [x_j] : \quad \frac{\partial u_i}{\partial x_j} \bigg/ \frac{\partial u_i}{\partial z_i} = \frac{\partial p}{\partial x_j} + R_j \cdot \frac{\partial H}{\partial x_j} \tag{1.6}$$

$$[R_j] : \quad u_i(x_j^r, s_i^r, \xi_j^r, z_i^r) \geq u_i(x_j^{-r}, s_i^{-r}, \xi_j^{-r} z_i^{-r}) \text{ for } R_j = r \in \{0, 1\}, \quad -r \neq r. \tag{1.7}$$

These first-order conditions are instructive. The left side of equation 1.6 is the household's marginal rate of substitution between consumption of the housing attribute and consumption of the outside good, while the right side is the price ratio with the cost of the outside good normalized to one. Note that the implicit price to which each household responds depends on whether the household sorts into a landlord-pay or tenant-pay regime. If in a tenant-pay regime, then $R_j = 1$, so choices of x_j (such as square feet) will affect the marginal cost of energy use. If in a landlord-pay regime, then $R_j = 0$, so choices of x_j no longer affect the marginal cost of energy use, eliminating the price incentive to conserve energy.[13] Equation 1.7 says that the household sorts into the payment regime that brings it the highest utility when it considers the optimal housing characteristics for each. The dependence of the first-order conditions on R_j confirms the intuition that households have different incentives for energy consumption under each payment regime. These conditions also illustrate the simultaneous choice of heating contract and other housing characteristics that necessitates studying housing and heating decisions in a joint framework.

How much do the energy-use incentive distortions matter? To answer this question, the difference in energy use under the current regime and under an alternative regime in which all households have tenant-pay contracts must be calculated, relying fundamentally on counterfactual housing and

[12] Linear separability abstracts away from risk aversion. Given regulated energy prices, price uncertainty plays a small role in this market. Empirical studies of farm tenancy (Allen and Lueck, 1992, 1999), timber contracts (Leffler and Rucker, 1991), groundwater contracts (Aggarwal, 2007), and even health insurance markets (Einav et al., 2013) reject risk aversion as a determinant of contract choice. Allen and Lueck (1995) review the literature of contract choice across several settings and conclude generally that risk aversion does not play a large role in contract choice.

[13] A similar condition exists describing the choice of the unobserved attribute ξ_j but is not the focus of the analysis.

heating scenarios that cannot be observed. While the housing and heating choices of households can be observed, household choices in a world with only landlord-pay or tenant-pay options cannot be observed. To identify the counterfactual energy outcomes and welfare implications, I estimate household utility functions with preferences revealed by household sorting and temperature choices. These utility functions can be used to estimate households' behavioral responses to policy changes and to calculate the welfare effects of moral hazard and choice of housing attributes.

1.3.2 Empirical model

Building upon Bajari and Kahn (2005) and Bajari and Benkard (2005), I specify a flexible utility function for households:

$$u_i(x_j, s_i, \xi_j, z_i) = \beta_{1,i} ln(x_j) - \frac{1}{2\beta_2}(s_i - s_i^b)^2 + \beta_{3,i} ln(\xi_j) + z_i. \tag{1.8}$$

Households have heterogeneous preferences for how warm they like to keep their homes as well as for consumption of housing attributes. Note that the utility parameters $\beta_{1,i}$ and $\beta_{3,i}$ for housing attributes in equation 1.8 vary by household. The term s_i^b is household i's heterogeneous bliss point temperature setting. Deviation from the bliss point temperature causes disutility at a rate that depends on parameter β_2 and the bliss point for household i.[14]

Monthly average heating cost for household i is the price of fuel P_e multiplied by the quantity of fuel used $Q_{e,i}$ and the binary regime indicator R_j. The quantity of fuel used to heat to a given indoor temperature setting is the difference in the temperature setting s_i from the outside temperature during winter T_j, scaled by the housing unit's efficiency level.[15] I assume that home attributes act as energy-efficiency complements (e.g., lower square footage and low temperature

[14]The distance norm is chosen for differentiability and simplicity. This parametric form is similar to that adopted in Einav et al. (2013). A more general functional form could be used that allows for stronger disutility from being too cold than too hot or vice versa, but this is not necessary given the focus on heating.

[15]To reflect heating-season temperatures, the outdoor temperature during winter T_j is interpreted as mean outdoor temperature conditional on being colder than 65°F. This is to reflect heating-season temperatures.

setting could have a multiplicative effect):

$$Q_{e,i} = (s_i - T_j)^{\sigma} \cdot x_j^{\gamma} \cdot \xi_j^{\kappa} \tag{1.9}$$

where σ, γ, and κ are fuel-intensity parameters. From this specification, the marginal heating cost of attribute x_j is

$$\frac{\partial H}{\partial x_j} = R_j \left(\gamma \cdot \frac{P_e \cdot Q_{e,i}}{x_j} \right), \tag{1.10}$$

while the marginal heating cost of the temperature setting s_i is

$$\frac{\partial H}{\partial s_i} = R_j \left(\sigma \cdot \frac{P_e \cdot Q_{e,i}}{s_i - T_j} \right). \tag{1.11}$$

In the model, I make the computational assumption that households respond to marginal heating costs averaged by year, census division, regime, and fuel type.[16]

In addition, I assume rents are a function of housing attributes x_j, fuel prices P_e, regime type R_j, and potentially unobserved housing or market attributes ν_{dmj}, according to the rent price function:

$$rent_{dmj} = \mu_{0,m} + \alpha_{1,d} x_j + \alpha_{2,d} P_{e,d} + \alpha_{3,d}(1 - R_j) P_{e,d} + \alpha_{4,d}(1 - R_j) + \mu_1 w_j + \nu_{dmj}, \tag{1.12}$$

where $rent_{dmj}$ is the price of unit j in census division d, and MSA m.[17] I allow coefficients on attributes to vary by census division and allow mean rents to vary by MSA according to $\mu_{0,m}$. The term w_j are controls in estimation used to account for unobserved unit heterogeneity.[18] Thus, $\alpha_{1,d}$ represents the marginal rent cost of attribute x_j in division d. The markup for having the landlord pay for utilities is the difference in energy price pass-through, $\alpha_{3,d} P_{e,d}$. A higher fuel price should be capitalized into the implicit rental markup for a landlord-pay housing unit. Different fuel prices

[16] While my specification provides estimates of completely heterogeneous marginal costs due to the functional form of heating costs, any given estimate for i is likely a poor estimate (Wooldridge, 2010). Taken on average over division, year, regime, and fuel type, the average marginal heating cost estimates are unbiased and consistent while preserving the variation due to fuel prices and average housing stock efficiency differences to identify heterogeneous effects.

[17] Census divisions are regional groups of states: New England, Mid Atlantic, East North Central, West North Central, South Atlantic/East South Central, West South Central, and Mountain/Pacific.

[18] For example, w_j includes indicators by unit in the fixed-effects specification or individual means in the correlated random effects specification.

should induce changes in markups for landlord-pay units with different fuel types over time. If the fuel price was equal to zero, it would not matter whether a housing unit's heating bill was paid by the landlord or the tenant. Thus, the interaction between fuel price and landlord-pay regime traces out the markup, while the landlord-pay term controls for unobserved differences correlated with regime type.

Just as before, households set an optimal heating rule and choose housing attributes. Households have unlimited wants and limited resources such that total spending on housing, energy services, and other goods must be less than or equal to income y_i. Under my funtional-form assumptions, the first-order conditions in equations 1.4, 1.6, and 1.7 yield

$$[s_i]: \quad \frac{-1}{\beta_2}(s_i^* - s_i^b) = R_{j^*}\left(\sigma \cdot \frac{P_{e,d} \cdot Q_{e,i}}{s_i - T_{j^*}}\right) \tag{1.13}$$

$$[x_j]: \quad \frac{\beta_{1,i}}{x_{j^*}} = \alpha_{1,d} + R_{j^*}\left(\gamma \cdot \frac{P_{e,d} \cdot Q_{e,i}}{x_{j^*}}\right) \tag{1.14}$$

$$[R_j]: \quad u_i(x_j^r, s_i^r, \xi_j^r, z_i^r) \geq u_i(x_j^{-r}, s_i^{-r}, \xi_j^{-r}, z_i^{-r}) \text{ for } R_j = r \in \{0,1\}, \quad -r \neq r. \tag{1.15}$$

These equations characterize the choices of home heat setting s_i, home attributes x_j, and energy contract regime R_j.

1.3.3 Parameter identification

Identification of parameter $\beta_{1,i}$ is straightforward. Solving for $\beta_{1,i}$ in equation 1.14 reveals that for the optimal choice $j^*(i)$:

$$\beta_{1,i} = x_{j^*}\left(\alpha_{1,d} + R_{j^*}\left(\gamma \cdot \frac{P_{e,d} \cdot Q_{e,i}}{x_{j^*}}\right)\right). \tag{1.16}$$

A household's choice of housing characteristic x_{j^*}, with estimates of the hedonic price $\alpha_{1,m}$ and marginal cost of heating reveal an estimate of parameter $\beta_{1,i}$. Estimation is discussed in more detail later, but these are relatively simple applications of a hedonic price regression and a "hedonic heat cost" regression. Note that the first-order conditions and the estimated parameters can be solved for the counterfactual choices of x_j under the two payment regimes.

Temperature preferences β_2 and s_i^b are more difficult to identify. Using the home heating rule in equation 1.13 for individuals in tenant-pay regimes ($R_j = 1$),:

$$s_i^* = s_i^b - R_{j^*}\left(\sigma \cdot \frac{P_{e,d} \cdot Q_{e,i}}{s_i - T_{j^*}}\right)\beta_2. \tag{1.17}$$

Thus, temperature settings by landlord-pay households ($R_j = 0$) reveal bliss-point temperature settings, so s_i^b is directly observed for landlord-pay households. Likewise, temperature settings by tenant-pay households reveal $s_i^* \neq s_i^b$ for tenant-pay households. With an estimate of the disutility parameter β_2 and the marginal cost of temperature setting, I can use equation 1.17 to back out the unobserved bliss point for tenant-pay households as well as the counterfactual optimal temperature setting under positive energy prices for current landlord-pay households. I use two complementary approaches to identify the key parameter β_2, both based on equation 1.17: I estimate β_2 using a revealed-preference approach in the RECS as well as design and implement a choice experiment via a nationally-representative survey to elicit changes in temperature setting in response to randomized variation in marginal heating costs.

First, I regress observed temperature setting on marginal costs, controlling for being in landlord-pay and observed demographics. The bliss point is a physiological parameter that depends on demographic characteristics D_i. I therefore model $s_i^b = S(D_i) + \delta(1 - R_{j^*}) + h_i$, where $S(D_i)$ is a function of demographics, $\delta(1 - R_{j^*})$ represents unobserved differences in bliss points for households which sort into landlord-pay units, and h_i represents household taste heterogeneity.[19] Using these assumptions and equation 1.13, the optimal heating choice can be rewritten as

$$s_i^* = \left[S(D_i) + \delta(1 - R_j) + h_{1,i}\right] - \left[\beta_2 MC(P_e, Q_{e,i}, s_i, T_{j^*}, R_{j^*})\right], \tag{1.18}$$

where $MC(\cdot)$ is the marginal cost of temperature setting:

$$MC(P_e, Q_{e,i}, s_i, T_{j^*}, R_{j^*}) = R_{j^*} \cdot \left(\sigma \cdot \frac{P_e \cdot Q_{e,i}}{s_i - T_{j^*}}\right). \tag{1.19}$$

[19] I later discuss the possibility that the bliss point temperature setting may be different in different housing units; e.g., a higher temperature may be preferred in an old and drafty unit relative to a new unit.

Equation 1.18 traces out each household's home heating rule. The first bracketed term represents the bliss point as a function of demographics, while the second bracketed term represents the response to the cost of home heating. Thus, a regression of temperature setting on household demographics and the marginal cost of temperature setting (estimated in the heating cost regression) provides an estimate of the mean disutility parameter β_2. As a side benefit, the coefficients on $(1 - R_j)$ and D have economically meaningful interpretations. The δ coefficient on $(1 - R_j)$ represents the mean difference in bliss point preference for households in landlord-pay regimes not explained by demographics. An estimate of $S(D_i)$ provides an estimate of each household's bliss point that is explained by observed demographic coefficients.

Second, to supplement this revealed-preference approach, I conduct a choice experiment that allows me to estimate equation 1.17 directly using exogenous variation in energy costs. To do so, I first elicit a subject's bliss point by asking at what temperature they would set their thermostats if heating was free. Next, I draw a low, medium, and high cost of changing the thermostat that is calibrated from the estimated heating costs. Finally, I regress temperature setting on marginal cost to get an estimate of β_2 that is free of endogeneity.

Thus, the model provides closed-form expressions for households' heterogeneous preferences $\beta_{1,i}$ and s_i^b. I use the estimated utility function parameters and heating costs to calculate partial-equilibrium counterfactual energy use scenarios for each household and characterize the relative effects of moral hazard, choice of home size, and sorting based on temperature preference. The model allows me to estimate the welfare costs of moral hazard and equilibrium sorting through deadweight loss and the external cost of emissions from fossil fuels.

1.3.4 Empirical model strengths and limitations

The functional-form assumptions on the utility function eliminate the need for several other assumptions common in the literature. First, the functional form allows the utility parameters to vary arbitrarily by household. Though I specify the parametric form of the utility function, the parameters are extremely flexible and general. Second, the random coefficients do not require a researcher-imposed error term as is common in most random coefficient utility models (Bajari and Kahn, 2005). Finally, some functional form assumptions are necessary to simulate realistic counterfactuals in any setting.

Another attractive feature of the estimation strategy is that it does not require any assumptions about supply-side market structure (Bajari and Kahn, 2005). Households react to prices set outside of their influence. These prices could arise in a perfectly competitive or monopolistic setting. The downside to this aspect of flexibility is that I cannot model general-equilibrium supply-side effects. For example, in a counterfactual simulation of removing the landlord-pay regime, landlords would likely respond by changing the level of efficiency in the existing housing stock. While housing attributes such as square footage are difficult to change, changes to appliances are a realistic response to a policy change. In addition, the hedonic prices of energy-efficiency features of housing units will adjust in general equilibrium. Computed counterfactuals therefore only reflect the response of households in a partial equilibrium.

In addition to these relaxed constraints, the estimated parameters have an intuitive reduced-form interpretation that lends credibility to the functional-form assumption. Equation 1.16 that identifies $\beta_{1,i}$ states that a household's preference for a housing attribute is revealed by the *amount* of a housing attribute x_j purchased as well as the *true cost*, $\left(\partial p/\partial x_j + R_j \cdot \partial H/\partial x_j\right)$, of the attribute. This makes sense: if household A purchases more square feet than household B at the same price, all else equal, then household A's preference for square feet $\beta_{1,A}$ is revealed to be greater than B's preference, $\beta_{1,B}$. Similarly, if household A faces higher cost per square foot than household B but chooses the same amount of square feet, it is revealed that $\beta_{1,A} > \beta_{1,B}$. Thus, the utility parameters incorporate information from variation in quantity chosen, the market price, and the cost to downstream heating.

The empirical identification of the utility parameters depends on the flexibility of the estimated gradients. To see this, consider estimating an *inflexible* price gradient by restricting the coefficient on square feet to be the same for all households in the United States. Call this estimate $\hat{p}_{sf}$, the marginal effect of square feet from a regression of price on housing attributes. Then consider two households:

household A chooses a 1,000-square-foot apartment with landlord-pay utilities in Los Angeles, California, and household B chooses a 1,000-square-foot-apartment with landlord-pay utilities in Raleigh, North Carolina. Using the framework developed above, $\hat{\beta}_{sf,A} = \hat{\beta}_{sf,B} = 1,000 \cdot \hat{p}_{sf}$. The issue is that the price of an additional square foot is not the same in Los Angeles and Raleigh—the household in Los Angeles has likely paid much more per square foot and has a stronger preference. The utility differences between the two households are not identified due to unobserved price differences between markets. In fact, the estimation of a single marginal price for square feet adds no more information beyond the level choice of square footage. In this example, accurate estimates of the marginal prices of square footage in each market identifies the difference between two households' utility parameters in different markets. The difference in quantities purchased identifies within-market variation in a household's utility parameters.

I proceed by estimating the model parameters used to construct the utility parameters above.

1.4 Estimation of model parameters

The key empirical objects needed to identify the model's parameters are the marginal impacts of thermostat setting and home size on heating cost as well as the hedonic price of landlord-pay utilities and home size. I estimate these objects using a heat cost regression and hedonic rent price regression. Using these estimates and individual home choices, I construct estimates of the preference for square feet $\beta_{1,i}$. I then use a regression-based approach to estimate the temperature disutility parameter β_2 and selection into landlord pay based on bliss point. I corroborate these estimates using data from a choice experiment. Armed with estimates of β_2, I am able to back out unobserved bliss points for current tenant-pay households. Finally, I use a LASSO regression to explore heterogeneity in bliss-point temperature.

1.4.1 Heat costs

This step estimates the causal effect of home attributes and temperature-setting behavior on home heating cost. How much more money does it cost to heat an additional square foot of home space, ceteris paribus? Given that fuel use is governed by the relationship in equation 1.9, the marginal heating cost of attribute x_j is $\left(\gamma \cdot (P_e \cdot Q_{e,i})/x_j\right)$ and the marginal heating cost of temperature setting is $\left(\sigma \cdot (P_e \cdot Q_{e,i})/(s_i - T_j)\right)$. I estimate γ and σ by log-linearizing equation 1.9:[20]

$$ln(Q_e) = \sigma ln(s_i - T_j) + \gamma ln(x_j) + \kappa ln(\xi_j), \tag{1.20}$$

where $s_i - T_j$ is the difference between the winter thermostat setting s_i and the average outdoor temperature conditional on being less than 65°F.[21] Of course, the chosen heating intensity $s_i - T_j$ is an endogenous function $q(\cdot)$ of the fuel price P_e, home efficiency due to attributes $ln(x_j)$, heating degree days HDD_j, demographics D_i, and unobserved tastes h_i:[22]

$$ln(s_i - T_j) = q(P_e, ln(x_j), HDD_j, D_i, h_i). \tag{1.21}$$

Equations 1.20 and 1.21 are a classic demand system that I estimate using two-stage least squares. Equation 1.20 is the structural equation of interest, in which $s_i - T_j$ is endogenous. Equation 1.21 points to a first-stage regression in which energy prices affect temperature setting but satisfy the standard exclusion restriction and can be used as an instrumental variable for temperature setting. The estimation sample is all tenant-pay renters in the RECS whose fuel use is obtained directly from

[20]For attributes with zero or negative values, I use the inverse hyperbolic sine transformation as an approximation to the natural log.

[21]The ideal but infeasible approach would be to calculate an individual measure of heating intensity for each household i that measures the cumulative daily difference between the indoor temperature setting and the outside temperature when the temperature outside is lower than the thermostat setting. Thus if $T^{ideal}_{j,d}$ is the daily outdoor temperature, the ideal measure would be $(1/365)\sum_{d=1}^{365} max\{0, s_i - T^{ideal}_{j,d}\}$. Because of data-privacy concerns, I do not observe each household's exact daily outdoor temperature over the year. Instead, I observe "65°F heating-degree days" for each household (HDD_j), which is the cumulative difference between 65°F and the outdoor temperature when it is less than 65°F outside (i.e., $HDD_j = \sum_{d=1}^{365} max\{0, 65 - (T^i_{j,d} deal | T^{ideal}_{j,d} < 65)\}$). Using the provided HDD_j, I back out the average outdoor temperature conditional on the outdoor temperature being less than 65°F and label this T_j. The constructed measure of heating intensity is the difference between the temperature setting s_i and the average outdoor temperature conditional on the outdoor temperature being less than 65°F.

[22]See footnote 21 for a definition of heating degree days, a commonly-used measurement of frequency and intensity of cold days per year.

the utility company.[23] Nearly all renters heat with either electricity or gas; therefore, I estimate heat regressions for electricity and gas only. The final sample of electricity-heated households includes a repeated cross section of 1,511 renters surveyed in 2001, 2005, and 2009. The final sample of gas-heated households includes a repeated cross section of 1,653 renters surveyed in 2001, 2005, and 2009.

Table A.4 displays the ordinary least squares and two-stage least squares estimates of equation 1.20. The two-stage least squares estimates of the energy-impacts of temperature setting σ and unit characteristics γ are precise. For gas-heated homes (with an average monthly bill of $59), a one percent increase in thermostat setting increases the heating bill by 0.49 percent, all else equal. A one percent increase in square feet increases the heating bill by 0.21 percent, all else equal. Similarly, for electricity-heated homes (with an average monthly bill of $91), a one percent increase in thermostat setting increases the heating bill by 0.2 percent, all else equal. A one percent increase in square feet increases the heating bill by 0.26 percent, all else equal.

Table A.5 displays the estimated marginal costs using the two-stage least squares estimates. Increasing the thermostat setting by one degree increases the heating bill by $3.57 on average for gas-heated homes. It costs an average of $1.01 per month to heat an additional 100 square feet for gas-heated homes. Increasing the thermostat setting by one degree increases the heating bill by $2.89 on average for electricity-heated homes. It costs $2.42 per month to heat an additional 100 square feet for electricity-heated homes. I use estimated marginal heating costs averaged by year, census division, and by heating fuel type to construct preferences and estimate welfare effects in the simulation model.

I also display the estimated savings from living in a larger apartment building. For gas- and electricity-heated homes, heating bills are lower in buildings with a larger number of units. This relationship between building size and energy savings is not due to the number of floors as argued by Borck and Brueckner (2017). These results suggest economies of scale in heating for large buildings with many units, independent of the number of floors.

[23] I exclude households whose fuel use is estimated. Note that fuel consumption and billing is not available for landlord-pay households.

1.4.2 Temperature disutility and bliss points

I first estimate the temperature disutility parameter (β_2) in a regression of temperature setting on estimated marginal cost using revealed-preference data from the RECS. My ordinary least squares estimates are precise but inconsistent, while my two-stage least squares estimates are consistent but imprecise. Thus, I use data from a choice experiment that I conducted to get a precise causal estimate of β_2 to corroborate the cross-sectional estimates. Finally, I explore potential heterogeneity in the disutility parameter using a LASSO approach. The point estimates of my revealed-preference approach are corroborated by the choice experiment data.

First, I directly estimate the temperature disutility parameter β_2 from the RECS using a specification suggested by the heating rule in equation 1.18. I regress temperature settings on demographics D_i, a dummy for being on landlord-pay heating $(1 - R_j)$, and the estimated marginal cost $\widehat{MC}$:

$$s_i = \psi_0 + \psi_1 D_i + \psi_2(1 - R_j) + \psi_3 \widehat{MC}_i + h_i, \tag{1.22}$$

where h_i is mean-zero heterogeneity. The marginal cost is estimated for household i in a first stage as described in the preceding section.[24] These predicted marginal cost terms are endogenous because they depend mechanically on temperature setting, so I use fuel prices as an instrument for temperature setting. Given the heating rule in equation 1.18, $\hat{\psi}_3$ is an estimate of the mean disutility parameter β_2, while $\hat{\psi}_2$ is the estimated difference in bliss-point temperature preferences for households in landlord and tenant pay due to factors not in D_i.

Table A.6 displays these estimates. The coefficient on the marginal-cost term is imprecise and sensitive to controls in the two-stage least squares specifications. Meanwhile, the ordinary least squares estimates are tightly estimated around -0.55. Surprisingly, I find that landlord-pay households have bliss points that are 1-3°F lower than tenant-pay households on average (67°F vs 69-70°F).

Next, to supplement imprecise estimates from the cross-sectional approaches, I conduct a

[24] Standard errors are bootstrapped and allow for sampling error in the estimation of the marginal cost term.

separate choice experiment to corroborate the estimates.[25] The steps are as follows: First, I elicit bliss point by asking subjects what temperature setting they would choose if there was no cost to setting the thermostat. Next, I draw a low, medium, and high marginal cost for temperature setting from three independent uniform distributions. The temperature marginal costs are calibrated from the estimates obtained from the RECS and range from $1 to $8 per month to change the thermostat by five degrees Fahrenheit. I then reveal each price to the respondents and ask at what temperature they would set their thermostat, tracing out each respondent's demand for energy services.

I conducted the choice experiment in March 2018 using the Qualtrics Survey Panel. The final sample includes 414 individuals drawn from a nationally representative sample who each completed three temperature setting choices; combined, these choices make 1,242 individual-choice observations.[26] Using the responses, I directly estimate equation 1.17 using the following specification:

$$s_{i,c} - s_i^b = a + bMC_{i,c} + \varepsilon_i, \tag{1.23}$$

where $s_{i,c}$ is respondent i's temperature setting for choice c and $MC_{i,c}$ is the randomly drawn marginal cost. Given this estimating equation, $\hat{b}$ is an estimate of β_2.

Table A.7 displays the ordinary least squares estimates, which indicate a β_2 of -0.8. The estimates are precise and are robust to controlling for individual fixed effects. These results are similar to those obtained from the RECS approach, lending confidence to the estimates above.

Across the well-powered estimation methods, the disutility parameter β_2 ranges from -0.5 to -1. For the final simulation, I use the fixed-effects estimate of $\hat{\beta}_2 = -0.79$ from the choice experiment and explore the sensitivity of the results to alternative parameter values of β_2 = -0.5 and -1. Figure A.12 displays the estimated bliss-point temperature settings for landlord and tenant-pay households assuming β_2=-0.79. Surprisingly, landlord-pay households have lower bliss points than tenant-pay households. Thus, I do not find strong evidence that households select into landlord-pay regimes

[25] As described in chapter 2.

[26] See the appendix for a more detailed description of sampling methodology, a description of how the final sample was chosen, and exact question text.

based on bliss-point temperature preference.[27]

1.4.2.1 Exploring heterogeneity with LASSO

Finally, I use a reduced-form double LASSO approach to flexibly explore heterogeneity in the bliss point while allowing for heterogeneity in the temperature disutility parameter β_2.[28] Here, I estimate an equation for temperature setting as a function of demographics D_i, housing attributes x_i including R_j, and heating fuel prices faced by tenant-payers $(R_j P_e)$.[29] I also include the interactions and squared terms of all these variables, resulting in 1,590 covariates. I estimate the equation

$$s_i = \Upsilon_0 + \Upsilon_1 x_j + \Upsilon_2 x_j^2 + \Upsilon_3 D_i + \Upsilon_4 D_i^2 + \Upsilon_5 x_j D_i + \Upsilon_6 (1 - R_j) \tag{1.24}$$

$$+ \Upsilon_7 R_j P_e + \Upsilon_8 R_j P_e^2 + \Upsilon_9 (R_j P_e) x_j + \Upsilon_{10} (R_j P_e) D_i + \eta_i \tag{1.25}$$

where η_i is unobserved mean-zero heterogeneity in preferences using the $L1$-norm LASSO penalization.[30] This approach flexibly traces out a predictive equation for temperature setting. I use this equation to predict bliss-point temperatures for all households when prices are equal to zero $\hat{s}_{i,lasso}|P_e = 0$ (landlord pay), and temperature settings for all households when prices are positive $\hat{s}_{i,lasso}|P_e > 0$ (tenant pay).

Figure A.7 displays the distribution of predicted bliss points using the LASSO algorithm. The figure implies that landlord- and tenant-pay households have nearly indistinguishable bliss points as predicted by observables. The distribution of estimated temperature disutility parameters $\beta_{2,i}$ in figure A.8 shows similar temperature disutilities by regime. Landlord-pay households are less responsive to marginal heating costs than tenant-pay households on average, though the distribution of parameters is very similar overall. Finally, figures A.9 and A.10 display the distributions of

[27] I explore alternative explanations to this finding in the appendix.

[28] The double LASSO uses a first-stage LASSO algorithm to select variables then performs ordinary least squares regression on the selected variables. This approach eliminates the parameter bias introduced by LASSO. See Belloni and Chernozhukov (2013).

[29] By using fuel prices rather than marginal cost in this regression I avoid the endogenous marginal cost terms.

[30] I choose the penalization weight λ using 10-fold cross validation.

predicted temperature settings under moral hazard and fully internalized prices for landlord-pay and tenant-pay households. Both groups of households display significant predicted response to paying for heating.

To demonstrate the improved performance of the LASSO over OLS, I include cross-validation performance metrics in table A.8 and plot the predicted temperature settings versus the true distribution in figure A.11. I compare the LASSO to ordinary least squares and find that LASSO significantly reduces mean-squared error and bias in cross validation. Despite these performance gains and the wide variety of included predictors, the root-mean-squared error of the best performing LASSO is still on the order of 5°F. I conclude that a large degree of heterogeneity in temperature preferences is idiosyncratic and difficult to predict.

1.4.3 Hedonic prices and preferences for attributes

The goal of this step is to recover the marginal prices of apartment unit attributes. For example, how much does it cost to rent an apartment with more square feet, ceteris paribus? Using the AHS data, I regress rental rate on unit characteristics while controlling for housing market and year effects. I allow for different marginal price coefficients by census division, which is the finest level of geography that can be used to link prices between the AHS and RECS.

I estimate the hedonic gradient, allowing coefficients to vary by census division:

$$rent_{dmj} = \mu_{0,m} + \alpha_{1,d}x_j + \alpha_{2,d}P_{e,d} + \alpha_{3,d}(1-R_j)P_{e,d} + \alpha_{4,d}(1-R_j) + \mu_1 w_j + \nu_{dmj}, \quad (1.26)$$

where $rent_{dmj}$ is the price of unit j in census division d, and MSA m. I allow coefficients on attributes to vary by census division and allow mean rents to vary by MSA according to $\mu_{0,m}$. The term x_j represents attributes of unit j, $(1-R_j)$ is a binary variable equal to one if the landlord pays for heat, $P_{e,d}$ is the heating fuel price, w_j are controls in estimation used to account for unobserved unit heterogeneity, and ν_{dmj} are unobserved attributes or market characteristics affecting rent.[31] Pass-through of regional electricity and natural gas prices into the rent identifies the landlord-pay

[31] I control for unit square feet, bedrooms, bathrooms, air conditioning type, in-unit laundry, number of units in the building, frequency of the rent payment, heating equipment type (e.g.,

markup $\Delta p/\Delta R$—similar to a strategy used in two recent papers by Myers (2017, 2018). If the fuel price was equal to zero, it would not matter whether a housing unit's heating bill was paid by the landlord or the tenant. However, a positive fuel price should be capitalized into the implicit rental markup for a landlord-pay housing unit. Further, variation in relative fuel prices should induce differential changes in markups for landlord-pay vs tenant-pay units relying on different fuel types over time. This intuition allows for the inclusion of a rich set of controls such as unit fixed effects, indicators for fuel type, and flexible vintage by year indicators. Because of the relative variation in fuel prices over time (depicted in figure A.3), the landlord-pay markup is still identified when including unit fixed effects. Intuitively, the identifying variation for the markup comes from exogenous changes of fuel prices that are capitalized into higher rents for landlord pay units. This strategy requires the assumption that rentals do not vary over time and space between payment regimes in unobserved ways that are correlated with changes in relative fuel prices over time. Figure A.3 displays the time-series variation in the relative fuel prices (in \$/MMBTU) used to identify the markup.

Table A.3 reports the estimates of the hedonic price of square feet and the estimated markups by census division. The specifications include a "vanilla" ordinary least squares estimation, a housing unit random-effects estimation, a housing unit fixed-effects estimation, and a housing unit correlated-random-effects estimation. The preferred correlated-random-effects specification uses the within-unit variation on fuel prices over time to trace out the markup for having the landlord pay for heating but uses across-unit variation in square feet to estimate the hedonic price of square feet. Figures A.5 and A.4 display the results from the preferred correlated-random-effects specification. I find that 100 additional square feet costs anywhere from \$5-\$20 extra per month to rent. These estimates are consistent across all specifications. The high estimates in New England

furnace, electric room heaters, etc), non-heating utilities that the landlord pays, whether the unit was too cold last year, a subjective rating of neighborhood quality, an indicator for gas heat, division interacted with fuel type, division indicators, vintage by year interactions, landlord-pay interacted with division, MSA, and year dummies interacted with the landlord-pay term. The w_j term includes indicators by unit int he fixed-effects specification or individual means in the correlated random effects specification.

and the Mountain/Pacific divisions are likely driven by expensive housing in New York City, San Francisco, and other such cities. The confidence intervals for the estimated markup for having heat included in the rent are wide and include zero for most census divisions. The point estimates suggest a markup as low as zero and as high as $100 per month.

I construct household preferences for square feet $\beta_{1,i}$ by substituting the estimated values for the hedonic price and average marginal heating cost into equation 1.16. Figure A.13 displays the distribution of estimated preferences for square feet. Landlord-pay households have significantly lower preferences for square feet than tenant-pay households on average and their choices are more sensitive to changes in the price per square foot.

1.5 Simulation results

Consider requiring all tenants to pay for their own heating bills. Landlord-pay households would respond by re-sorting into different-sized housing units and changing their temperature settings. I use the estimated utility parameters for temperature setting and housing unit size to estimate household responses to such a change. Table A.9 contains the results of the simulation. Column (1) uses the preferred estimate of the temperature disutility parameter $\beta_2 = -0.79$ and columns (2) and (3) display results for alternative values of β_2. On average, landlord-pay households re-sort into housing units that are 140 square feet smaller and reduce their thermostat settings by 4 degrees Fahrenheit. Overall, these changes translate into lower energy expenditures of $218 per household per year or a 25 percent decrease relative to observed energy bills.[32] While 140 square feet seems like a large decrease in home size, it only accounts for $30 per household per year. The moral hazard effect dominates, accounting for a full $188 of lower expenditures per household per year.

What do these results imply for the welfare losses associated with the landlord pay regime? I calculate the deadweight loss from moral hazard and the sorting effect by considering not just additional expenditures but also each household's benefit from consuming additional energy services and square feet when the landlord pays. In addition, I calculate the additional carbon dioxide emis-

[32]This change is relative to the average bill for tenant-pay units because the actual heating bill is not observed for landlord-pay households.

sions due to increased demand for fossil fuels.[33] The total calculated welfare loss is the deadweight loss plus the external cost of the carbon dioxide emissions under landlord pay relative to a scenario in which all tenants pay for their own heating and sort into different-sized housing units.[34]

Table A.9 breaks down the yearly welfare cost per household and for the entire United States. I calculate an average welfare loss of \$154 dollars per household per year; external emissions costs account for \$43.68 of this total. Using the RECS survey weights, I calculate the annual welfare loss across all landlord-pay households in the United States. I find that the total welfare loss is \$839 million per year. Of this amount, \$238 million is due to the external cost of carbon dioxide and \$601 million is due to the deadweight loss. I also calculate the welare costs based on the high and low range of disutility parameter (β_2) estimates and find that the total welfare loss ranges from \$682 million to \$1.26 billion per year.

The model estimates significant regional heterogeneity in welfare impacts. Both the East North-Central and West North-Central regions have the lowest total deadweight loss impacts at \$117 and \$70 per household per year. The Mountain/Pacific and South Atlantic/East South Central regions have the highest total deadweight loss impacts at \$203 per household per year. Regional differences in fuel prices and the number of carbon-intensive electricity-heated units drive this heterogeneity.

1.6 Conclusions

When the landlord pays for heating, the renter does not face any price incentive for energy conservation—a clear form of moral hazard. Landlord-pay households thus choose higher thermo-

[33] I assume a \$40/ton external cost of carbon dioxide. For natural gas, I assume that natural gas combustion releases 53.12 kilograms of carbon dioxide per 1000 cubic feet (EIA, 2016). For electricity, I calculate a simple carbon emission intensity for each census division by dividing total electricity production by total carbon emissions from electricity production.

[34] In terms of the model, the total monthly welfare loss to society from moral hazard in landlord pay household i is $\int_{s_i|R_j=1}^{s_i^b} \frac{\partial H}{\partial s_i} + \beta_2 \left(s - s_i^b \right) + \mathfrak{C}_s(e,m) ds$ where $\mathfrak{C}_s(e)$ is the marginal carbon cost of temperature setting for fuel type e in census division m. The total monthly welfare loss to society from sorting on attribute x_j is $\int_{x_j|R_j=1}^{x_j|R_j=0} \left(\frac{\partial H}{\partial x_j} + \frac{\partial P}{\partial x_j} \right) - \frac{\beta_{1,i}}{x} + \mathfrak{C}_x(e,m) dx$ where $\mathfrak{C}_s(e)$ is the marginal carbon cost of square feet for fuel type e in census division m.

stat settings than if they paid for heating. In addition, when the landlord pays for heating, the renter does not pay the implicit energy cost of of square feet and other home attributes. Landlord-pay households thus choose larger housing units than if they paid for heating. Landlord-pay residential energy contracts create perverse energy-use incentives for moral hazard and equilibrium sorting that if eliminated could provide private savings and reductions in fossil-fuel emissions.

In this paper, I build a model in which households jointly choose home attributes and temperature setting. Households sort into housing units by selecting the heating contract and the size of the unit, considering the difference in heating cost. After selecting a home, the household chooses the indoor temperature setting. If in a tenant-pay unit, the household faces a tradeoff between energy services and other consumption. If in a landlord-pay unit, the household ignores heating costs and heats to its temperature bliss point. The model allows me to estimate the welfare impacts of choosing a housing unit that is too large and setting the thermostat too high.

Using observed rent prices, home attributes, energy bills, temperature settings, and a choice experiment, I estimate key parameters in the model. I find that landlord-pay households heat 4°F higher than if they were made to pay their own heating bills. In addition, these households choose housing units 140 square feet larger than if they were made to pay their own bills. These distortions result in welfare losses of about $154 per household per year. Aggregated across all US landlord-pay households, these results imply $839 million in welfare losses per year from landlord-pay contracts, including $238 million due to the external cost of carbon emissions.

The landlord's role in contract choice is still not well-understood. Future work should explore the landlord's role in selecting housing attributes and choosing to bundle utility payments into the rent. I find suggestive evidence for economies of scale in home heating: for gas-heated units, units in larger apartment buildings are significantly less expensive to heat. Units in larger apartment buildings are also significantly more likely to have landlord-pay utilities. Landlord-pay units may have more energy-efficient appliances and insulation than tenant-pay units (Myers, 2018; Gillingham et al., 2012). The results here and in the literature suggest that landlords choose to pay for heating when it is profitable. The landlord's joint decision has not received much attention in the current literature and would shed light on the other side of the market.

CHAPTER 2

THE REMARKABLY INELASTIC DEMAND FOR HOME ENERGY SERVICES: A CHOICE EXPERIMENT

2.1 Introduction

The law of demand states that as the price of a good increases, the quantity of the good a consumer demands will decrease. When applied to environmentally damaging goods, the law of demand suggests that prices are a powerful tool for conservation. Since the seminal work of Pigou (1920), environmental economists have studied how prices of polluting goods can be increased to reduce consumption and account for the external costs of pollution. Energy is the quintessential polluting good. Energy production typically requires burning fossil-fuel inputs such as oil, natural gas, or coal.[1] Combustion of fossil fuels produces costly local pollutants (e.g., particulate matter, nitrogen dioxide, and sulfur dioxide) and global pollutants (e.g., carbon dioxide). At the same time, the energy services derived from fossil fuels provide profound benefits to human well-being (e.g. lighting, heating, and other electronic technologies). Understanding demand for energy is important for creating institutions that properly incorporate environmental costs into decision making.[2]

Central to this story is the assumption that individuals respond to changing energy prices and trade off the benefit from energy consumption with the cost of energy consumption. Recent empirical work challenges this assumption. One study of electricity-use data finds that 44 percent of studied households did not respond to prices at all (Reiss and White, 2005). An analysis of natural gas billing data finds that both low-income and high-income households do not respond to prices (Auffhammer and Rubin, 2018). An experiment in Sweden where renters were switched from landlord-pay to tenant-pay electricity shows that while average electricity consumption decreased by

[1]In 2017, of all useful thermal output in the United States, 37 percent came from non-fossil fuels. In the electric power sector, only 9.6 percent of thermal output came from non-fossil fuels (EIA, 2018).

[2]Energy demand is also a key concern for electricity and natural gas utilities. For example, if the supply and demand of electricity is not balanced at all times during a day there will be a blackout.

24 percent, two-thirds of the reduction came from just 20 percent of the studied households (Elinder et al., 2017). Price responsiveness is low for most energy users but high for a select group. Why do some individuals fail to respond to prices while other individuals cut energy consumption drastically when prices are high? There are two main potential explanations for low price responsiveness: behavioral heterogeneity and preference heterogeneity. Either non-price-responsive individuals are uninformed or face high costs to monitor energy prices, or these individuals know prices and rationally choose not to respond because of high valuation of energy services.[3]

This paper tests whether individual heating-choice behavior is consistent with households having full information about the cost of energy. I conduct and analyze a stated-choice experiment with a nationally representative sample of individuals in which participants make choices about how high to set their thermostat during the winter when told the hypothetical cost of doing so. In this experimental setting, energy costs are easy to understand, costless to monitor, and salient. The experimental environment is clean of any potential confounding factors such as unobserved energy efficiency, thermostat or meter placement, and attrition bias which makes it difficult to interpret results from field experiments. The results from the choice experiment serve as a fully informed benchmark to compare to real temperature-setting behavior. If hypothetical temperature-setting behavior and real temperature-setting behavior match, this provides evidence for heterogeneous preferences as the primary driver for energy-use heterogeneity. If experimental temperature-setting behavior differs from real temperature-setting behavior, this is evidence pointing to price salience, informational barriers, and adjustment costs as determinants of energy-use heterogeneity.

From the experiment, I find that under perfect experimental conditions, 50 percent of individuals do not change the thermostat at any treatment cost. On average, a 100 percent increase in the cost of heating the home one degree Fahrenheit warmer induces a 0.31-0.97 degree Fahrenheit reduction

[3]The literature suggests a number of behavioral responses or informational barriers for energy use. For example, Ito (2014) find evidence that energy users respond to average rather than marginal prices. Jessoe and Rapson (2014) argue that consumers do not know prices or face a high cost of determining energy prices. Allcott and Rogers (2014) find that social comparisons impact energy use and observe behavior consistent with short attention spans. Finally, Allcott and Taubinsky (2015) argue that consumers do not pay attention to energy prices when choosing light bulbs.

in the winter heating level. In addition, I find that reported actual household temperature-setting behavior is consistent with realistic beliefs about the cost of heating. This evidence suggests that preference heterogeneity plays a large role in driving empirical observations of inelastic and heterogeneous energy demand. I analyze the heterogeneous price responses and find that individuals with higher temperature preferences are more price responsive. This is evidence that increasing block pricing of emissions can reduce emissions while minimizing the regressive properties of emissions fees.

2.2 Experimental procedure

Experiment participants comprise a nationally representative sample of US individuals drawn from the online Qualtrics Panel. I eliminated respondents if they failed Qualtrics speeding checks, failed attention check questions placed in the survey (e.g. “Agree or disagree: I breathe more than once per day."), if they do not use heat at home in the winter, or if they provided poor-quality responses (e.g. gibberish in free response boxes). The final sample includes 414 individuals. The experiment took place in early March 2018, the end of winter for most of the United States; thus, respondents completed the survey after making heating decisions for several months.

First, I elicit each individual's temperature preference baseline by asking what temperature they would choose if there was no cost to heat:[4]

> *Imagine that you do not have to pay for heating your home during the winter. In this situation, what temperature setting (degrees F) would you choose when you are at home?*

[4]The science and engineering literatures argue that temperature preference is determined by physiological characteristics such as age (Taylor et al., 1995; Schellen et al., 2010), sex (Kingma and van Marken Lichtenbelt, 2015; Karjalainen, 2012, 2007; Fanger, 1970; Parsons, 2002; Cena and de Dear, 2001; Muzi et al., 1998; Pellerin and Candas, 2003; Griefahn and Künemund, 2001; Nakano et al., 2002; Nagashima et al., 2002), diet (Ringsdorrf Jr. and Cheraskin, 1982), and previous exposure (Young, 2010). There is some evidence that temperature preferences of men and women differ by country (Beshir and Ramsey, 1981; Karjalainen, 2007; Indraganti and Rao, 2010) and that individuals may be able to consciously alter the body's internal response to temperature (Kox et al., 2014).

This baseline temperature preference with no price can be thought of as a *bliss point temperature preference* for heating. Next, respondents see an example of how much money would be saved for a variety of thermostat settings.

I draw a low, medium, and high marginal cost from three indepdendent uniform distributions spanning \$1 to \$8 per month for a five degree Fahrenheit change when they are home.[5] Respondents see a price and are asked to input their chosen temperature setting. For example,

> *Choice #3: Imagine increasing your thermostat by one degree Fahrenheit will increase your monthly heating bill by \$1.60 (or changing your thermostat by five degrees Fahrenheit will increase your heating bill by \$8).*
>
> *When a one degree change in temperature costs \$1.60 per month, what temperature setting would you choose?*
>
> *Remember that you said you would set your thermostat to 70 degrees Fahrenheit if you weren't paying for heating.*

Respondents input their chosen temperature into a text-response box. After completing the experiment, respondents supply their demographic information. Qualtrics compensates each respondent a small sum after participating successfully. A copy of the survey instrument is available for viewing in the supplementary dissertation files.

2.3 Data

Table C.1 displays summary statistics from the experimental sample after cleaning the data. The sample is higher income, older, more female, and more white than the nationally representative quotas used to recruit subjects. Figure C.1 displays kernel density plots of participants' bliss point temperature preferences and actual temperature settings. The distribution of bliss point temperatures appears to have a higher mean and similar variance to the distribution of actual

[5] The first price is a random draw from a U(1,2.67) distribution, the second price is a random draw from a U(2.67,5.33) distribution, and the third price is a random draw from a U(5.33,8) distribution.

temperature settings. A Kolmogorov-Smirnov test of equivalence of distributions easily rejects the null hypothesis that the distributions are the same.

Half (54 percent) of participants report that they set their actual thermostats equal to their bliss point temperature preference. Figure C.2 shows a plot of bliss point vs actual temperature setting. 6.5 percent of participants report that the thermostat is higher than the bliss point, perhaps because they did not understand the question or because they are not in control of the thermostat. For experimental temperature settings, 50 percent of participants continue to choose their bliss point temperature setting at the highest cost level (including 73 percent of individuals who set their actual home temperature equal to their bliss point).

2.4 Estimating demand for heating

The experiment provides points on each respondent's temperature demand curve. The most intuitive measure of an individual's temperature response to a change in the price of an additional degree is the *semi-elasticity*, or the degree change in the thermostat for a percent change in price. For each experiment choice $c \in \{1, 2, 3\}$, an individual i sees a price $price_{i,c}$ to change the thermostat and chooses a temperature setting $temp_{i,c}$. Thus, for a household i, I model the choice of temperature setting as some function $f(\cdot)$ of the price:

$$temp_{i,c} = f(price_{i,c}) + \varepsilon_i, \tag{2.1}$$

where ε_i is individual heterogeneity that is orthogonal to $price_{i,c}$. The semi-elasticity is

$$(\partial temp/\partial price) \cdot price_{i,c}. \tag{2.2}$$

First, I pool the sample and estimate the mean semi-elasticity using ordinary-least-squares and fixed-effects estimation. I estimate the following equation on the pooled temperature choices:

$$temp_{i,c} = \alpha + \beta \ln(price_{i,c}) + \epsilon_i.$$

Taking the derivative with respect to the price variable and solving for β reveals that $\beta = (\partial temp/\partial price) \cdot price_{i,c}$. Thus with this functional form, the estimate $\hat{\beta}$ serves as an estimate of the average semi-elasticity.

Next, I estimate each individual's unique semi-elasticity for temperature setting by calculating the arc semi-elasticity directly and using a regression adjustment method. Following the intuition developed in Allen and Lerner (1934), I calculate the arc semi-elasticity between each pair of points on each participant's temperature demand curve. The arc semi-elasticity between two choices tracing out the demand curve c and $c-1$ is

$$\text{Arc semi-elasticity}_{c,c-1} = \frac{\Delta_{c,c-1} temp_i}{\%\Delta_{c,c-1} price_i}, \tag{2.3}$$

where $\Delta_{c,c-1} temp_i = temp_{i,c} - temp_{i,c-1}$, the difference in temperature settings chosen by the participant, and

$$\%\Delta_{c,c-1} price_i = \left(\frac{price_{i,c} - price_{i,c-1}}{\frac{price_{i,c} + price_{i,c+1}}{2}} \right), \tag{2.4}$$

the percentage difference in researcher-assigned energy price. One benefit of this approach is that the arc semi-elasticity can utilize information from the bliss point choice (i.e., when price is zero), while the regression-based approaches cannot due to the log of zero being undefined. In addition, this approach provides a heterogeneous and non-parametric measure of price responsiveness.

Finally, I measure each participant's semi-elasticity using a regression-adjustment approach. For each participant i, I estimate the following equation separately with ordinary least squares:

$$temp_{i,c} = \alpha_i + \beta_i \ln(price_{i,c}) + \epsilon_i. \tag{2.5}$$

The estimate $\hat{\beta}_i$ is an estimate of each individual's mean semi-elasticity over individual i's experimental choices c. This method provides heterogeneous semi-elasticities but does not incorporate information provided from the bliss point choice.

Table C.2 displays the estimated average semi-elasticities using all four methods. I bootstrap the 95 percent confidence intervals of the averages using 1,000 replications and re-sampling at the participant level. The estimates imply that for a 100 percent increase in the cost of heating, an individual reduces the thermostat setting by 0.31-0.97 degrees Fahrenheit. This small average response is due to the large number of price-insensitive participants and hides significant heterogeneity, which I analyze in the following section.

2.5 Heterogeneity analysis

Figure C.3 displays a histogram of participants' arc semi-elasticities, and figure C.4 displays a histogram of participants' regression-adjustment semi-elasticities. The distributions display a similar bunching of individuals completely unresponsive to prices with a significant portion of more price-responsive individuals in the tail.

The distribution of elasticities is characteristic of those found in other energy settings. Reiss and White (2005) estimate a similarly skewed distribution of annual elasticities for electricity use with a mass of relatively price-insensitive households and a fat tail of more elastic households. They also find that low-income households are more elastic and conclude that households with space heating are significantly more elastic than other households. The experiment here shows that the skewed distribution can be generated without the energy-efficiency responses included in a yearly elasticity.

I explore what drives heterogeneity in temperature response by regressing the arc semi-elasticities on standardized vectors of the average price on the arc $\overline{price}_{i,c,c-1}$ and participant demographics $demographics_i$. I use a Tobit maximum-likelihood estimation to account for the clustering at zero in the dependent variable. Thus, denoting $Z(\cdot)$ as the function that transforms a sample draw of a random variable into its z-score, I estimate the equation

$$\text{Arc semi-elasticity}_{i,c,c-1} = a + b\mathrm{Z}(\overline{price}_{i,c,c-1}) + d\mathrm{Z}(demographics_i) + e_{i,c,c-1} \tag{2.6}$$

using maximum likelihood, treating all non-negative arc semi-elasticities as a corner solution. Standardization allows the marginal effects of the regression to be easily compared. The marginal effects from this estimation are interpreted as the change in arc semi-elasticity for a one-standard-deviation change in the predictor variable holding the other predictor variables constant.[6]

I explain heterogeneity as a function of bliss point temperature preference, average monthly heating bill, income, age, household size, number of children living at home, gender, race, ur-

[6] The marginal effect I estimate is the "unconditional" average partial effect $\frac{\partial \mathbb{E}(\text{Arc semi-elasticity}|Z(x))}{\partial z(x_j)}$ where x is a matrix of predictor variables and x_j is a single predictor variable using the results provided in (Wooldridge, 2010).

ban/rural status, education, and political party. Figure C.5 plots the estimated marginal effects with the 95 percent confidence intervals bootstrapped using 1,000 replications with repeated sampling at the participant level. Most strikingly, individuals with a one-standard-deviation-higher bliss point temperature have on average a -.42 higher arc semi-elasticty (i.e. are more elastic), all else equal. Higher-income and higher-education respondents are less elastic on average, all else equal, although the confidence intervals for the education marginal effects include zero. Participants living in urban areas are more responsive to price changes. Older participants are less elastic, with a one-standard-deviation increase in age corresponding with a .23 lower arc semi-elasticity, all else equal. The marginal effects estimates of average heating bill, participant gender, race, number of children, and household size have confidence intervals that contain zero.

Political party is not a strong determinant of elasticity, with Republicans, Democrats, and Independents having statistically indistinguishable elasticity measures when controlling for other covariates. In two papers, Costa and Kahn estimate heterogeneous energy use patterns by political ideology. First, Costa and Kahn (2013a) show that total household electricity use is lower for politically progressive households. Second, Costa and Kahn (2013b) finds that politically progressive homeowners are more responsive to non-price nudges. The experiment in this paper measures a different dimension of energy use, but nonetheless the lack of heterogeneity by political group is surprising. It is possible that in the literature, total energy use and ownership of energy-efficient appliances are correlated with progressive energy-efficiency policies and thus reflect these policies rather than individual behavior. In this estimation, I include many controls that are correlated with ideology and whose influence may be spuriously attributed to ideology (e.g., urban or rural).

In this experiment, age plays a large role in determining elasticity whereas sex does not. While the science and engineering literatures focus on measuring differences in temperature preference and sensitivity, differences in *behavior* are often ignored. While a group of people may on average be able to detect a difference in temperature in a laboratory more readily, this does not translate necessarily to differences in thermostat setting behavior. Indeed, I find here that men and women do not respond to prices differently after other characteristics have been controlled for despite numerous findings that women prefer higher temperatures than men.[7] The findings in the literature may reflect how temperature decisions are made in settings that affect multiple individuals with heterogeneous temperature preferences or other barriers to adjusting the thermostat.[8]

2.6 Conclusion

The experiment reproduces energy-use heterogeneity distributions comparable to those seen in actual energy-use data. Half of participants report that they set their actual thermostats equal to their bliss point temperature preference when at home. Of these individuals, 70 percent were similarly unresponsive to the cost of heating in the experiment. This is evidence that for these 70 percent of individuals, there is some perceived negative preference for deviating from their temperature bliss point in excess of the savings that they could have made in the experiment. These participants' behavior is consistent with a rational zero response to the cost of heating at the relevant price level. Under perfect-information conditions, energy-use behavior features significant heterogeneity and unresponsiveness.

There are two main empirical findings in the paper. First, a majority of individuals set their thermostats consistent with having perfect information. It is not likely that every individual knows the exact cost-per-degree change on the thermostat, but over time most people have adjusted their behavior based on feedback from energy bills. Second, more than half of all individuals are completely unresponsive to prices. People simply do not like to be cold. The cost of heating is low enough to take heating for granted, but it is likely that even if the cost of heating was to dramatically increase (perhaps due to a pollution fee), behavior would respond very little. Inelastic demand for energy services does not mean that prices are ineffective; instead, it means that the benefits from energy services are high. As long as the inelasticity does not arise from an artificial barrier such as false information about the energy cost savings, individuals will make the proper tradeoff between

[7] See Karjalainen (2012) for a review of this literature.

[8] For example, Kingma and van Marken Lichtenbelt (2015) discuss temperature demand in shared office buildings and Karjalainen (2007) finds that women are less likely to change the thermostat settings than men are.

costs and benefits from energy use when facing prices that reflect the full external costs of energy use.

These results imply that the gains from home energy efficiency are likely high. If energy-use behavior is fixed for many individuals, energy efficiency savings are large and will not be cannibalized by a rebound effect. It is not clear whether households optimally adopt energy-efficiency upgrades (i.e., whether there is an “energy-efficiency gap"), but a recent review of the literature did not find much evidence that individuals systematically fail to adopt energy efficiency (Gillingham and Palmer, 2014).

One important implication of these findings is that increasing block pricing can be used to reduce energy use without large incidence for a bulk of inelastic energy users.[9] The largest determinant of elasticity in the experiment was having a high bliss point temperature preference, implying that large energy users are more price responsive. By increasing the price of energy for large energy users who are most price-responsive, a regulator or regulated energy provider can reduce load (and corresponding emissions) without increasing payments from inelastic users. For example, a carbon tax with a zero-price carbon allowance may not sacrifice efficiency gains but may reduce the regressivity of the policy.

[9] Increasing block pricing charges a higher marginal cost per unit of energy for consumption of units of energy over a threshold. It essentially provides users with an allowance of cheap energy each billing period before having to spend more on additional energy consumption.

CHAPTER 3

NEGAWATTS VS. MEGAWATTS: DEMAND RESPONSE IN WHOLESALE ELECTRICITY MARKETS (WITH KATHERINE WAGNER)

3.1 Introduction

Between 2007 and 2017 in the United States, retail consumers spent \$377 billion per year on electricity—about 2.3% of annual GDP (EIA, 2018; BEA, 2019). Before a utility sells electricity to a retail consumer, the utility purchases electricity from generators on real-time wholesale markets. Despite electricity market liberalization in the 1990s, the typical retail rate is regulated and does not vary with the real-time marginal cost of electricity; thus, consumers do not receive price signals of scarcity when electricity is expensive during peak hours. During peak hours, the wholesale price of electricity often exceeds the retail price, and utility companies operate at a loss during these hours to keep the lights on. To address variability, regulators and utilities use a number of “demand-side management" tools to induce reductions in electricity demand when the price is high. One of the most controversial of these tools is wholesale demand response.

Wholesale demand response (hereafter “demand response") is a Federal Energy Regulatory Commission (FERC) program that allows an end user or middleman to sell a reduction in electricity use on the wholesale electricity market as if that reduction was generation. In a typical wholesale electricity market, renewable and fossil-fuel generators make bids to supply megawatts of electricity in an hour as long as the wholesale price exceeds the generator’s reserve price. The market operator then accepts the cheapest generation required to meet demand in the hour and pays all suppliers the bid of the marginal generator. With demand response, an end user or middleman can participate in the wholesale market by submitting a bid to supply “negawatts" or reductions in electricity use in an hour. If the demand response bidder’s reserve price is less than the market-clearing wholesale price, the bid is accepted, and the bidder must curb electricity use by the accepted amount or pay a fine. The utility company pays the bidder the wholesale price multiplied by the amount of reductions supplied. The logic is that it is cheaper to pay for reduced electricity demand than it is to produce these marginal units during peak periods when energy production would be inefficient and expensive. The demand response compensation scheme was the subject of a 2016 Supreme Court case in which the Court ultimately ruled it did not have the expertise to question FERC's demand response compensation rate, and that FERC was justified in maintaining a demand response program "that no one...disputes will curb prices and enhance reliability in the wholesale market" (Christiansen, 2016). Certainly no one disputes the benefits of a program that properly passes wholesale prices to electricity consumers; however, institutional design details dictate the incentives and the ultimate costs and benefits of the program.

In this paper, we analyze electricity consumers' incentives to contract with middlemen and reduce electricity use through the demand response program. By mandating that utilities compensate reductions at the wholesale price, FERC allows real-time prices to pass through to retail consumers during peak hours; however, the economic incentives created by the program are not well understood. The Supreme Court's ruling upheld FERC Order 745, which mandates that utilities compensate reductions in electricity at the wholesale electricity price. Thus, for an electricity consumer participating in demand response, the opportunity cost of consuming electricity includes *both* the retail rate and the wholesale price of electricity. However, most consumers do not directly participate in demand response. Most demand response is contracted through a middleman aggregator called a "curtailment service provider."[1] The aggregator contracts with any number of electricity consumers and can be any entity participating in the wholesale electricity market, including utilities or third-party firms. Previous work on demand response ignores the role of the aggregator and relies on informal conceptions of how demand response operates.[2]

[1] For example, in the first 9 months of 2018 over 80 percent of demand response came from four curtailment service providers (Monitoring Analytics, 2018).

[2] Eryilmaz et al. (2017) provide some suggestive evidence that wholesalers may be more price-responsive during demand response periods but they ignore that wholesalers who do not reduce their demand during demand response periods pay a higher effective price because they forgo compensation for each unit of electricity purchased. Walawalkar et al. (2008) build a model of demand response that does not include the retail rate and ignores the effect of aggregators. O'Connel et al. (2014) review the state of demand response and note a possible effect of aggregators, but do

To understand the incentives created by demand response, we build a theoretical model of the PJM electricity market with demand response. Under minimal assumptions, we show that incentives for electricity use are not surplus-maximizing under the current demand-response rules. In addition, we show that direct participation, third-party aggregation, and utility aggregation create different (and inefficient) electricity-use outcomes. When end users participate directly in demand response, the combined opportunity cost of the retail rate and wholesale price creates an incentive to provide too many reductions. When there is an aggregator, only some of the incentive to reduce is passed through to the end user. The degree of pass through depends on the demand elasticity and which entity is the aggregator. Because the utility loses the retail rate for each unit of electricity not sold, the utility has a lower incentive to seek demand-response reductions than a third-party aggregator who does not forgo retail revenue.

Next, we compare the relative size of the economic surplus losses under different regimes. The effect on surplus depends on the slope of the electricity demand curve and the difference between the retail electricity rate and the wholesale electricity rate. Using a linear supply and demand framework, we show that when demand is less elastic and the retail rate is low, third-party aggregators perform better than the other second-best alternatives. When demand is more elastic and the retail rate is higher, utility aggregators perform better than the other second-best alternatives. We also show that direct participation can reduce surplus relative to the no-demand-response baseline.

Finally, we suggest that to maximize surplus, only utilities should participate in demand response. Because the utility is the residual claimant of surplus from load reductions, the utility is the best steward of demand response. In addition, regulators should recognize that demand response is a channel through which to exercise market power. A profit-maximizing utility does not have an incentive to pass on the surplus-maximizing compensation for reductions.

---

not dwell deeply on the subject.

3.2 Institutional background

Wholesale demand response is a load reduction program that is connected to the wholesale market and receives a price formally tied to the current wholesale price of electricity. In the PJM service area, wholesale demand response is designated as either “emergency" or “economic." Emergency demand response is actively called upon by the market operator (PJM) during a peak demand event. Economic demand response allows electricity users to submit reductions from baselines in electricity use on the wholesale electricity market as generation (PJM, 2014a). Retail demand response is a load reduction program operated by a utility that is not formally connected to the wholesale market and is similar to the interruptible load studied by Caves et al. (1988, 1992). In this paper, we focus on wholesale economic demand response as implemented in PJM. Further discussion of demand response refers to the wholesale “economic" variety.

Without demand response, utilities buy electricity from generators on a regional real-time wholesale market and re-sell it to electricity consumers at a regulated rate. On the wholesale side of the market, renewable- and fossil-fueled generators submit bids to generate megawatts of electricity at a minimum reserve price in each hour. PJM orders these bids from lowest to highest cost, creating an electricity wholesale market supply curve. On the retail side of the market, end users pay a regulated price for electricity determined by state regulatory agencies. These retail rates are typically disconnected from the real-time marginal cost of generating electricity and instead reflect regulatory desire for low and stable retail prices with a reliable grid. The utility is compelled by regulation to purchase electricity on the wholesale market sufficient to supply any level of demand at the fixed retail rate. The wholesale price of electricity the utility pays is equal to the bid of the last unit of generation required to meet demand. The utility is required to meet this demand even when the wholesale price of electricity is higher than the retail rate. Thus, end users do not have any extra incentive to conserve electricity during peak hours and subsequently demand more electricity than the surplus-maximizing level.

Demand response is designed to pass on real-time wholesale price fluctuations to end users. The idea is to subsidize reductions in electricity use by allowing end users to participate directly in the wholesale market. This raises the end user's per-unit opportunity cost of consuming electricity without raising retail rates—raising retail rates is outside of FERC's jurisdiction (Christiansen, 2016). Via demand response, end users participate in the wholesale market by submitting reductions in electricity use ("negawatts") as bids on the wholesale market. PJM treats a demand response bid exactly like an energy supplier's bid to generate electricity. The demand-response bid consists of an hourly marginal cost curve of reductions. If the market-clearing price exceeds the price of the bid for the demand-response reduction, PJM accepts the bid, and the bidder must provide the reductions within 20 percent of the amount accepted. The utility then must compensate accepted bids at the locational marginal price of electricity (the wholesale price plus delivery cost) plus shutdown costs. If the reduction is outside of the 20 percent window, the bidder forfeits any shutdown compensation, is charged a "balancing operating reserve charge," and receives the locational marginal price of electricity multiplied by the actual reductions. The utility pays for the demand-response reductions and no longer purchases actual generation on the wholesale market (PJM, 2013).

PJM determines what qualifies as a reduction by calculating the baseline electricity demand ahead of time. There is an initial certification process during which the demand response provider submits hourly meter data from 30 of the past 60 days. Once certified, the baseline load adjusts for each demand response event based on the electricity use within the past five days and on the day of the event (PJM, 2014b). Baseline manipulation is possible both during the certification process and during the same day adjustment (Kema, 2011; Chen and Kleit, 2016).

The market operator only accepts demand response bids when the locational marginal price of electricity exceeds a "net benefits threshold" set by PJM. The net benefits threshold is determined monthly by PJM and is meant to represent the price above which reductions in electricity use is worth soliciting (Rich, 2017). Between April 2012 and January 2013, the wholesale price exceeded the net benefits threshold for 75 percent of all hours (PJM, 2013).

Reductions must come from electricity end-users with an interval meter on the PJM market; however, any individual or entity is eligible to submit bids of reduction on the PJM wholesale market (PJM, 2014b). PJM refers to bidders as "curtailment service providers." In theory, curtailment service providers can be the electricity end-users themselves, utilities, or any third party entity. In practice, many third party firms have emerged as aggregators of demand response reductions and mainly handle establishing baseline load and submitting bids (McAnany, 2018). While PJM compensates the curtailment service providers at the locational marginal price, contracts between the curtailment service provider and the electricity consumer can take any form (PJM, 2017). We examine these contracts formally.

3.3 A model of demand response incentives

Here we develop an analytical framework to study the differing incentives of demand response aggregators.

Consider an electricity utility purchasing electricity generation on a competitive wholesale electricity market and selling electricity to consumers at a retail opportunity cost of p_r. In the wholesale electricity market, generators supply electricity at marginal cost and get paid the cost of the last accepted bid to generate electricity. Denote the wholesale market-clearing price as p_w and the wholesale market-clearing price without demand response as p_w^0. In the retail electricity market, the electricity price is constrained to a fixed retail rate r by regulation, so without demand response, the opportunity cost of electricity is the retail rate: $p_r = r$. During peak hours without demand response, the wholesale price is higher than the retail price $p_w^0 > r$, so the utility operates at a loss during those hours.[3]

We denote $D(p)$ the aggregate electricity demand at opportunity cost p in an hour and $S(p)$ the electricity supplied at opportunity cost p in an hour.[4] Let $MB(q)$ be the corresponding aggregate marginal benefit from consumption of q units of electricity and $MC(q)$ be the corresponding aggregate marginal cost of generating q units of electricity. The retail price cannot fluctuate and thus the inelastic aggregate hourly electricity demand is $I = D(r)$, but the wholesale price paid by the utility increases to induce supply that meets demand. Regulation compels the utility to

[3] In practice, state regulators allow the utility to recover costs through a flat fee per customer.

[4] Aggregate demand is downward sloping in p: $\frac{\partial D}{\partial p} > 0$. Supply is upward sloping in p: $\frac{\partial S}{\partial p} > 0$.

supply any level of quantity demanded at the regulated price. Thus, the equilibrium wholesale price equates supply with demand $D(p_r) = S(p_w)$, and the market-clearing wholesale price without demand response p_w^0 is found by solving $D(r) = S(p_w^0)$ for p_w^0.

The total surplus in the electricity market is the total benefit to consumers minus the total cost of generation at the quantity demanded at the retail opportunity cost $D(p_r)$. Define the total surplus function $W(D(p_r))$ in the electricity market as total benefits minus total generation costs of electricity generation:

$$W(D(p_r)) = \int_0^{D(p_r)} MB(\phi) - MC(\phi) d\phi. \tag{3.1}$$

Total surplus is maximized when the marginal benefit of electricity consumption equals the marginal cost, which is achieved when retail consumers face the same opportunity cost as wholesale suppliers. Thus, the surplus-maximizing electricity consumption is at $D(p_w)$, and the retail opportunity cost equals the market-clearing wholesale price of electricity $p_r^* = p_w$. When the retail opportunity cost is set equal to the wholesale price, the market-clearing wholesale price equals the retail opportunity cost.

We examine demand response incentives for electricity consumers and aggregators during peak hours.

First consider an electricity consumer i consuming electricity in a peak hour.[5] In the hour, customer i has an underlying demand curve $D_i(p)$ with corresponding marginal benefit $MB_i(q)$.[6] Because the retail price is fixed, inelastic demand for electricity is I_i, which is determined by the underlying demand at the fixed retail price: $I_i = D_i(r)$. When the wholesale price exceeds the retail price, the consumer uses too much electricity relative to the surplus-maximizing level $q_i^* = D_i(p_w^*)$.

Now consider a demand-response program designed to incentivize reductions in electricity consumption. The demand response program consists of a compensation payment to the consumer $c > 0$ per unit of reductions from baseline electricity consumption, i.e. $I_i - q_i$.[7] The cost

[5]The electricity consumer could be an industrial, commercial, or residential customer.

[6]Aggregate demand is the summation of individual demand curves in the market: $D(p) = \sum_i D_i(p)$.

[7]We describe how c is set and who pays later.

of consuming q units of electricity in the hour is then $rq_i - c(I_i - q_i)$.[8] The corresponding marginal opportunity cost of electricity consumption is then the retail rate plus the demand-response payment, i.e. $r + c$. When choosing how much electricity to consume, the electricity consumer sets the marginal benefit from electricity consumption equal to the marginal opportunity cost and consumes $q_i^{DR} = D_i(r + c)$. Thus, a demand-response incentive equal to $c^* = p^w - r$ leads the consumer to respond to the wholesale price and demand the surplus-maximizing amount of electricity $q_i^{DR}(c^*) = q_i^* = D_i(p_w^*)$.

The demand-response program lowers the equilibrium wholesale price of electricity. With demand response, the aggregate demand for electricity is $D(r + c)$. The equilibrium wholesale price of electricity with demand response p_w^{DR} clears the market so that all demand is supplied: $D(r + c) = S(p_w^{DR})$. Because demand is downward sloping and supply is upward sloping, the equilibrium wholesale price with demand response is lower than without, i.e., $p_w^{DR} < p_w^0$. With the optimal demand-response compensation rate, $c^* = p_w - r$, the equilibrium wholesale price satisfies $D(p_w^*) = S(p_w^*)$ and maximizes total surplus.

Figure D.1 plots the aggregate electricity supply and demand curves. Maximum surplus is at point q^*. Area 1 corresponds to the surplus loss in the no-demand-response baseline where consumers demand I units of electricity and the market-clearing wholesale cost is p_w^0. Area 2 corresponds to the surplus loss for a demand response compensation level $c > p_w - r$ that reduces demand past the surplus-maximizing level.

In practice, FERC requires that the utility pays p_w per unit of accepted demand-response reductions. Further, the compensation rate passed through to the final electricity consumer depends on whether the consumer participates directly in the demand response market or through an aggregator. In the next sections, we analyze how aggregators set the compensation rate c and the differing incentives for electricity use when consumers participate directly, through a third-party aggregator, or through a utility aggregator. We analyze the impact of each case on the equilibrium wholesale price of electricity. Finally, we compare the ultimate effect on total surplus in each case.

[8]This only holds for $q_i \leq I_i$, but we omit this condition from the remaining notation for simplicity.

3.3.1 Direct retail consumer participation

Given demand response, a retail consumer i can sell reductions in electricity use from baseline consumption I_i and directly receive compensation $c_1 = p_w$ per unit of reductions. Thus for a direct participator, the marginal opportunity cost of electricity use is the retail rate plus the wholesale price $p_r = r + p_w$. The consumer sets marginal cost equal to marginal benefit and consumes $q_1 = D_i(r + p_w) < D_i(p_w)$. The consumer reduces electricity consumption too much because the opportunity cost for consumption is artificially too high.

If all consumers participate directly, the aggregate electricity demand is downward sloping in the wholesale price: $D(r + p_w)$. The equilibrium wholesale price then satisfies $D(r + p_w^{direct}) = S(p_w^{direct})$. With direct participation, the equilibrium wholesale price of electricity drops too far relative to the optimal wholesale price and too little electricity is consumed.

Figure D.2 is a graph with the wholesale price on the vertical axis and quantity of electricity on the horizontal axis. The optimal amount of electricity consumption is at q^* when $D(p_w) = S(p_w)$. As before, in the baseline without demand response, the surplus loss is equal to area 1. Under direct participation in demand response, consumers always have an opportunity cost of the retail rate r. Thus, the demand curve is shifted downward to $D(p_w + r)$.[9] The intersection where $D(p_w + r) = S(p_w)$ is the equilibrium outcome. The equilibrium wholesale price decreases, but the total demand-response compensation increases. Area 2 corresponds to the surplus loss under direct participation in demand response.

The relative effect on total surplus depends on the slope of the electricity supply and aggregate demand curves. To generate useful comparisons between the cases, we calculate total surplus using linear specifications of supply and demand.[10] In particular, let aggregate demand be $D(p) = \alpha - \beta p$ and let the marginal cost of electricity supply be $MC = \gamma q$.[11] With the linear demand and supply

[9] When the wholesale price is zero, $D(p_w + r) = D(r) = I$, which can be seen on the graph.

[10] This assumption is meant to generate comparisons. To the extent that electricity supply and demand are locally linear this may be a good approximation to reality, but one should not read into this assumption as empirically motivated.

[11] One may also specify $MC = \delta + \gamma q$, but for ease of explanation we normalize $\delta = 0$ (or equivalently the demand intercept includes the supply intercept).

specifications, total surplus W is

$$W = \frac{1}{2}(\alpha - \beta p_r)\left(\frac{\alpha}{\beta} + p_r - p_w\right), \tag{3.2}$$

where $p_w = \gamma(\alpha - \beta p_r)$ and $p_r = r + c$. Taking the first order conditions for maximization and solving for c shows that the surplus-maximizing demand response compensation is $c^*_{linear} = (\alpha\gamma)/(1+\beta\gamma) - r$ with corresponding equilibrium wholesale price $p^*_{w,linear} = (\alpha\gamma)/(1+\beta\gamma)$. In other words, the optimum compensation eliminates the gap between the efficient wholesale price and retail opportunity cost. When there is no demand response, the equilibrium wholesale price is $p^0_{w,linear} = \gamma(\alpha - \beta r)$.

We return to the case of direct participation by end consumers. With linear aggregate demand, the equilibrium wholesale price is $p_w^{direct} = (\gamma\alpha - \beta\gamma r)/(1+\beta\gamma)$. Thus, the equilibrium demand-response compensation is $c_{1,linear} = (\gamma\alpha - \beta\gamma r)/(1+\beta\gamma)$. The equilibrium compensation is larger than the optimum by $r/(\beta\gamma+1)$. Thus, the distortion is increasing in the retail rate r and decreasing in β and γ. We compare the relative surplus effects in section 3.3.4.

3.3.2 Third-party aggregator

Assume that a third-party aggregator (called a curtailment service provider) contracts with retail consumers for demand-response reductions. The third-party aggregator acts as a middleman, receiving p_w per unit of reductions sold. The third-party aggregator offers consumers a payment of c_2 per unit of electricity reduced from the baseline and each consumer uses $q_{2,i} = D_i(c_2 + r)$ units of electricity.[12] We assume that a monopolist aggregator takes the wholesale price as fixed but can exert market power because it does not compete with other aggregators when setting the demand-response compensation.[13] The aggregator sets c_2 to maximize the arbitrage profits:

$$c_2 = \arg\max_c (p_w - c)(I - D(r + c)). \tag{3.3}$$

[12]We assume that there is a single aggregator for this exercise and therefore has market power in the arbitrage.

[13]If there is perfect competition in the aggregation arbitrage market and zero cost of aggregation services, then the third party aggregator case is equivalent to the direct participation case.

The arbitrage profits are less than or equal to zero when $c \geq p_w$ and when $c \leq 0$. Thus, $0 < c_2 < p_w$ and electricity demanded is $q_2 < D(r + c_2) < I$. The first-order condition for equation 3.3 determines the profit-maximizing c_2:

$$-p_w D'(r + c) = I - D(r + c) - cD'(r + c). \tag{3.4}$$

The left-hand side of equation 3.4 is the marginal arbitrage revenue and the right-hand side is the marginal arbitrage cost. Thus, the aggregator sets marginal revenue equal to marginal cost when choosing c_2. Finally, the market-clearing wholesale price p_w^{3rd} satisfies the condition $D(r + c_2) = S(p_w^{3rd})$. This is a downward shift in the demand curve, so the equilibrium wholesale price decreases relative to the zero-demand response case.

The relative effect on total surplus depends on the slope of the electricity supply and aggregate demand curves. We again consider a linear aggregate demand of $D(p) = \alpha - \beta p$ and linear marginal cost of electricity supply of $MC = \gamma q$ in order to compare the size of the distortion. Directly applying the result from the first-order condition in equation 3.3, the aggregator sets $c_2 = \frac{p_w}{2}$. Combining the first-order condition with the equilibrium condition, the equilibrium wholesale price is $p_{w,linear}^{3rd} = (\gamma\alpha - \beta\gamma r)/((\beta\gamma)/2 + 1)$. Thus, the equilibrium demand response compensation is $c_{2,linear} = 0.5(\gamma\alpha - \beta\gamma r)/((\beta\gamma)/2 + 1)$. We compare the relative surplus effects in section 3.3.4.

3.3.3 Utility aggregator

Assume that the utility contracts with consumers for demand response reductions. The utility saves $p_w - r$ for each unit of electricity reduced from the baseline (because it no longer has to purchase it on the wholesale market but also no longer sells it on the retail market) and offers a payment to electricity consumers of c_3 per unit of reductions. The utility sets c_3 to maximize savings from participating:[14]

$$c_3 = \arg\max_c (p_w - r - c)(I - D(r + c)) \tag{3.5}$$

$$\text{s.t. } p_w = MC(D(r + c)). \tag{3.6}$$

Note that the utility is a large enough player to know how the reductions will influence the marginal cost of electricity generation. The savings function in equation 3.5 makes it clear that $0 < c_3 < p_w - r$. If $c_3 \leq 0$, electricity consumption will not fall and the utility does not save money. If $c_3 = p_w - r$, the utility passes all the savings on to the electricity consumer, and if $c_3 > p_w - r$, the utility pays out more than the marginal benefit for reductions. Thus, the utility pays too little for reductions relative to the surplus maximizing level. The first-order condition for optimization determines the profit maximizing c_3:

$$\begin{aligned}(r - p_w)D'(r + c) + (I - D(r + c_3))\, MC'(D(r + c_3))D'(r + c_3)\\ = (I - D(r + c_3)) - cD'(r + c)\end{aligned} \tag{3.7}$$

The left-hand side of equation 3.7 is the marginal revenue and the right-hand side is the marginal cost. Thus, the utility sets marginal revenue equal to marginal cost when choosing c_3.

The relative effect on total surplus depends on the slope of the electricity supply and demand curves. We again consider a linear aggregate demand of $D(p) = \alpha - \beta p$ and linear marginal cost of electricity supply of $MC = \gamma q$ in order to compare the size of the distortion. Applying the result from the first-order condition in 3.7, the utility sets $c_3 = 0.5((\gamma\alpha)/(1+\beta\gamma) - r)$. Combining this with the equilibrium condition, the equilibrium wholesale price is $p^{3rd}_{w,linear} = 0.5\gamma(\alpha - \beta r + (\alpha)/(1+\beta\gamma))$. In the next section, we compare the relative surplus effects.

3.3.4 Comparing total surplus in the linear case

We now compare the total surplus from each regime in the linear case. The equilibrium deadweight loss from retail opportunity cost of electricity $p_r = r + c$ is

$$DWL = \frac{1}{2}|D(p_r) - D(p_w^*)||p_w - p_r|, \tag{3.8}$$

[14]This is equivalent to minimizing costs.

where $p_w = MC(D(p_r))$. Substituting the equilibrium compensation levels and wholesale prices from each of the cases, we find that the deadweight losses are:

$$DWL^{NoDR}_{linear} = \left(\frac{\beta}{2+2\beta\gamma}\right)(\alpha\gamma - r - \beta\gamma r)^2, \tag{3.9}$$

$$DWL^{direct}_{linear} = \left(\frac{\beta}{2+2\beta\gamma}\right)(r)^2, \tag{3.10}$$

$$DWL^{3rd}_{linear} = \left(\frac{\beta}{2+2\beta\gamma}\right)\left(\frac{1}{4}\right)\left(\frac{\alpha\gamma - 2r - 2\beta\gamma r}{1+\frac{\beta\gamma}{2}}\right)^2, \tag{3.11}$$

$$DWL^{Utility}_{linear} = \left(\frac{\beta}{2+2\beta\gamma}\right)\left(\frac{1}{4}\right)(\alpha\gamma - r - \beta\gamma r)^2. \tag{3.12}$$

These deadweight losses are difficult to compare analytically, but are easy to compare computationally. First, we consider 100 retail rates $r \in [40, 100]$, the equivalent to a typical PJM retail rate between 40 and 100 dollars per megawatt hour. We set $\alpha = 120,000$ and use 100 values of $\beta \in [100, 350]$. With these parameters, the demand elasticity at the retail price ranges from -0.43 to -0.03, covering a wide range of electricity demand elasticities estimated in the literature.[15] Without demand response, total load ranges from 85,000 to 116,000 megawatt hours, roughly matching the above-average (or peak) real-time loads experienced in PJM from 2010-2017 (Bowring, 2018). For supply, we set $\gamma = 0.0011$, leading to wholesale prices ranging from 77 to 104 dollars per megawatt hour which matches the above-average (or peak) wholesale electricity prices seen in PJM (Bowring, 2018).

First, we plot the deadweight losses by regime type as we vary the retail rate for a fixed level of the slope of demand, β. Figure D.3 shows how the deadweight loss varies with r. As the retail rate increases, the gap between the wholesale price and the retail price of electricity shrinks. Thus, the no-demand-response case becomes less damaging as the retail rate increases. In the case of direct participation by electricity consumers, as the retail rate increases, the total opportunity cost of electricity overshoots the optimum level by a larger amount and increases losses. Similarly, when

[15] The most commonly estimated elasticities are yearly and typically range from -0.1 to -0.5 on average (Deryugina et al., 2019; Reiss and White, 2005, e.g.). Long-run elasticities are typically larger. Hourly elasticities are smaller, for example around -0.16 in Jessoe and Rapson (2014). The chosen values are meant to span the feasible set of elasticities.

a third-party aggregator sets demand-response contracts, the total opportunity cost of electricity for the consumer is too high and increases too quickly. The deadweight loss increases at a slower rate than in the direct-participation case because the third-party aggregator reduces the demand-response compensation as the retail price increases. Finally, when the utility is the aggregator, the deadweight loss decreases as the retail rate increases. Similar to the third-party-aggregator case, the utility reduces the demand-response compensation as the retail rate increases.

Next, we plot the deadweight losses by regime type as we vary the demand elasticity parameter β. Figure D.4 shows how the deadweight loss varies with β. In the no-demand-response baseline, deadweight losses are increasing and then decreasing in elasticity. When demand is perfectly inelastic, there is no deadweight loss. When electricity consumers directly participate in demand response, the losses are increasing in elasticity because the consumers respond more strongly to the artificially-inflated opportunity cost of electricity consumption. This is much the same story when there is a third-party aggregator—increased price responsiveness increases the distortion. When the utility is the aggregator, the deadweight loss exhibits a quadratic shape similar to the no-demand-response baseline.

Finally, we examine the relative sizes of the deadweight losses from the regimes over all the parameter values tested. Figure D.5 is a set of heat maps of the deadweight losses from the regimes over the tested values of the retail rate and demand parameter. Each horizontal axis is the range of retail rates and each vertical axis is the range of demand parameters β. The darker colors are higher deadweight losses and the lighter colors are smaller deadweight losses. The gray area on the heat maps is the set of elasticity/retail price combinations at which the wholesale price is less than the retail rate.[16]

No regime dominates. Over the range of tested values, the third-party and utility aggregators outperform the no-demand-response and direct-participation cases. The utility-aggregator case has the lowest average deadweight loss overall, though when elasticity is low and retail rates are low, the utility aggregator performs better.

[16]This occurs when the retail rate is high and demand is more elastic. When demand is more elastic, the wholesale price is lower and thus closer to the retail price.

The no-demand-response baseline heat map illustrates that without demand response, the deadweight loss is larger when the retail rate is low and elasticity is high. When retail rates are low, electricity consumers do not have a strong incentive to conserve electricity, and thus surplus is lost from too much electricity consumed. However, the surplus loss from a too-low retail rate is smaller if the consumer is less price elastic because behavior under the surplus-maximizing price is not that different.

When consumers participate directly in demand response, the deadweight loss is larger when the retail rate is high and elasticity is high. Because directly-participating consumers are overcompensated for reducing consumption by the retail rate, the size of this distortion directly increases in the retail rate. When consumers are more elastic, they respond more to the improper incentive and reduce their electricity below the surplus-maximizing level even more.

The third-party aggregator's heat map shows that the deadweight loss is larger when the retail rate is high and elasticity is high. The third-party aggregator's pricing incentives face two partially-offsetting distortions. First, the third-party aggregator's marginal revenue per unit of reduction includes the full wholesale price rather than the difference between the wholesale price and the retail price, creating an incentive to over-provide reductions. Thus when the retail price is higher, the marginal revenue distortion is larger. Second, the third-party aggregator's market power creates an incentive to reduce reductions by not passing through the full demand-response incentive. This is easily seen from the aggregator's profit function in equation 3.3. If the firm compensates consumers at a full p_W per unit of reductions, profit goes to zero.

Finally, the utility aggregator's heat map shows that the deadweight loss is larger when the retail rate is small and customers are more elastic. In this case, the utility's marginal revenue is $p_W - r$ and the remaining distortion is because the savings from demand response are zero if the utility passes on the full compensation to consumers. The utility as aggregator is the only case in the current demand-response formulation in which surplus is guaranteed to increase from the zero-demand-response baseline (easily seen by noting the deadweight loss in this case is always one-fourth the deadweight loss without demand response in equations 3.9 and 3.12).

3.4 Can FERC improve demand response?

The analysis above shows that current wholesale demand-response institutions fail to incentivize surplus-maximizing electricity consumption. In fact, unless the utility is the aggregator, offering demand response is not even guaranteed to increase electricity market surplus. The FERC-mandated demand-response program requires utilities to compensate third-party and direct-participating consumers for reductions in load. Thus, while wholesale demand response passes real-time wholesale price variation on to electricity consumers, the institutional design creates poor incentives for electricity consumption.

Even without a FERC mandate, utilities have an incentive to enroll customers in programs such as retail demand response or other time-varying price programs. Utilities supply demand at fixed retail prices, so the utility is the residual claimant for any programs that increase the opportunity cost of electricity consumption during peak hours. The only remaining inefficiency is due to market power, which is currently dealt with through rate-of-return regulation or retail choice on a state-by-state basis.

This begs the question: why even have wholesale demand response? Proponents of wholesale demand response often argue that the program provides a path for behind-the-meter renewable generation to participate in wholesale markets (O'Connel et al., 2014). While perhaps true, allowing negawatts to be sold as megawatts at the same time reduces efficiency. If FERC wants to provide a path for distributed generation, it should pursue this goal directly through the wholesale market rather than through a convoluted demand-response program.

3.5 Conclusion

In this paper, we demonstrate that the aggregator is a potential source of inefficiency in FERC's wholesale demand response program. Using a theoretical model, we show that direct participants, third-party aggregators, and utility aggregators have different incentives to provide reductions in electricity use. In the current structure of demand response, direct particpants over-provide reductions because the opportunity cost includes the retail rate plus the wholesale-price incentive.

Third-party aggregators similarly have an incentive to over-provide reductions because of the compensation scheme but also have an incentive to under-provide reductions by exercising market power. Utility aggregators have an incentive to under-provide reductions if the utility can exert market power.

We argue that the wholesale demand-response institution is inefficient and superfluous to retail demand-response programs. The compensation scheme does pass on some real-time-wholesale-price variation to electricity consumers, but the pass-through does not create efficient electricity-use incentives and can even reduce economic surplus. The primary beneficiary of demand-response pricing is the utility, which can adopt retail demand-response or time-varying pricing programs separately from the wholesale market. Thus, we recommend that FERC end its wholesale demand-response program in favor of retail demand-response programs.

APPENDICES

APPENDIX A

FIGURES AND TABLES FOR CHAPTER 1

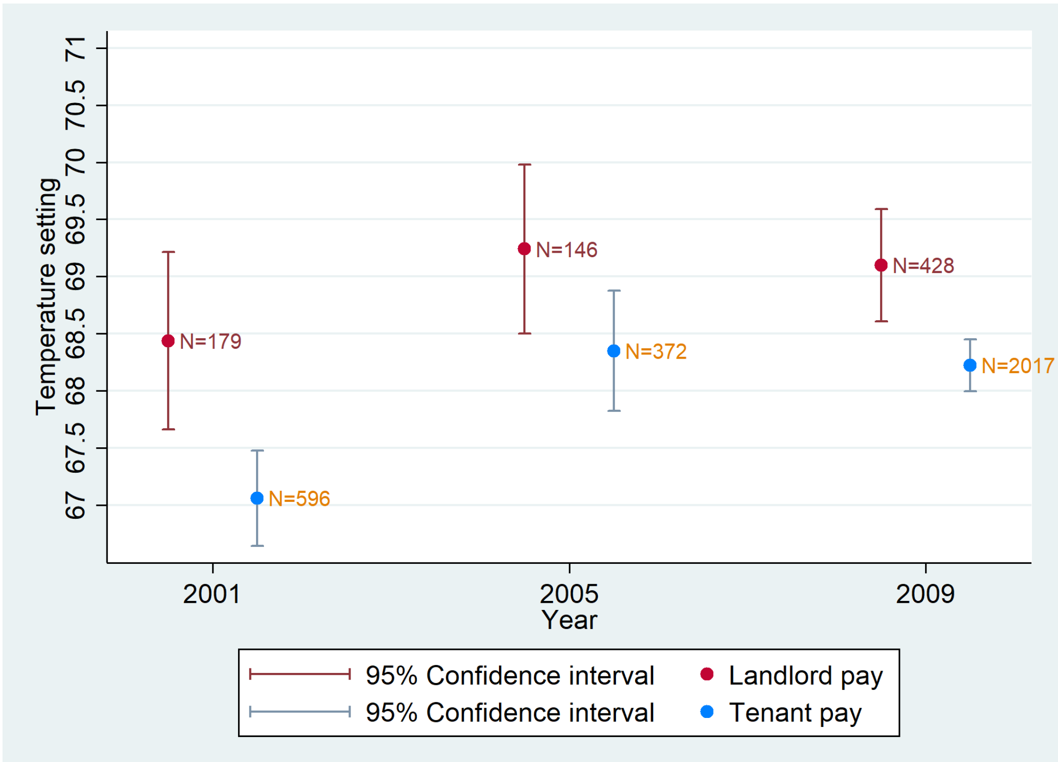

Figure A.1: Average temperature setting by year compared between landlord-pay and tenant-pay regimes. The tenant-pay temperature settings statistically vary between sample years, possibly due to changes in energy prices and resulting behaviors. This figure does not include households that turn off the heat completely: 28 (3.1%) landlord-pay households and 244 (6.8%) tenant-pay households.

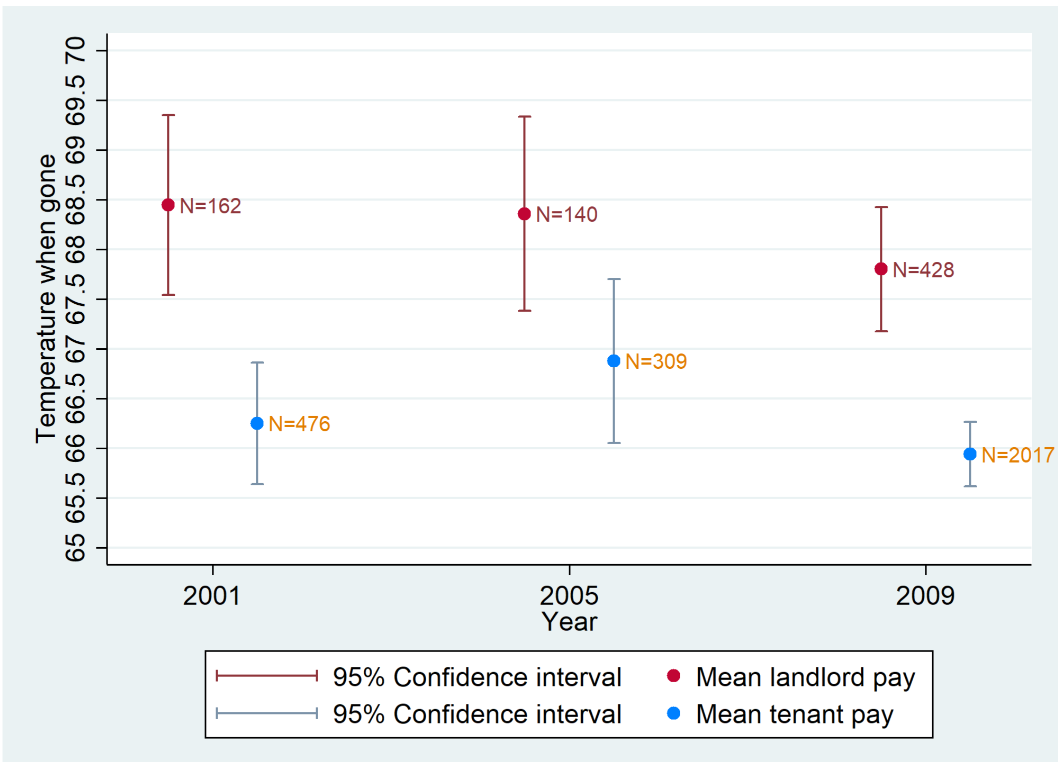

Figure A.2: Average temperature setting when out of the house between landlord-pay and tenant-pay regimes. This figure does not include households that turn off the heat completely: 28 (3.1%) landlord-pay housholds and 244 (6.8%) tenant-pay households.

Figure A.3: Time series variation in the average ratio of electricity to natural gas price per BTU.

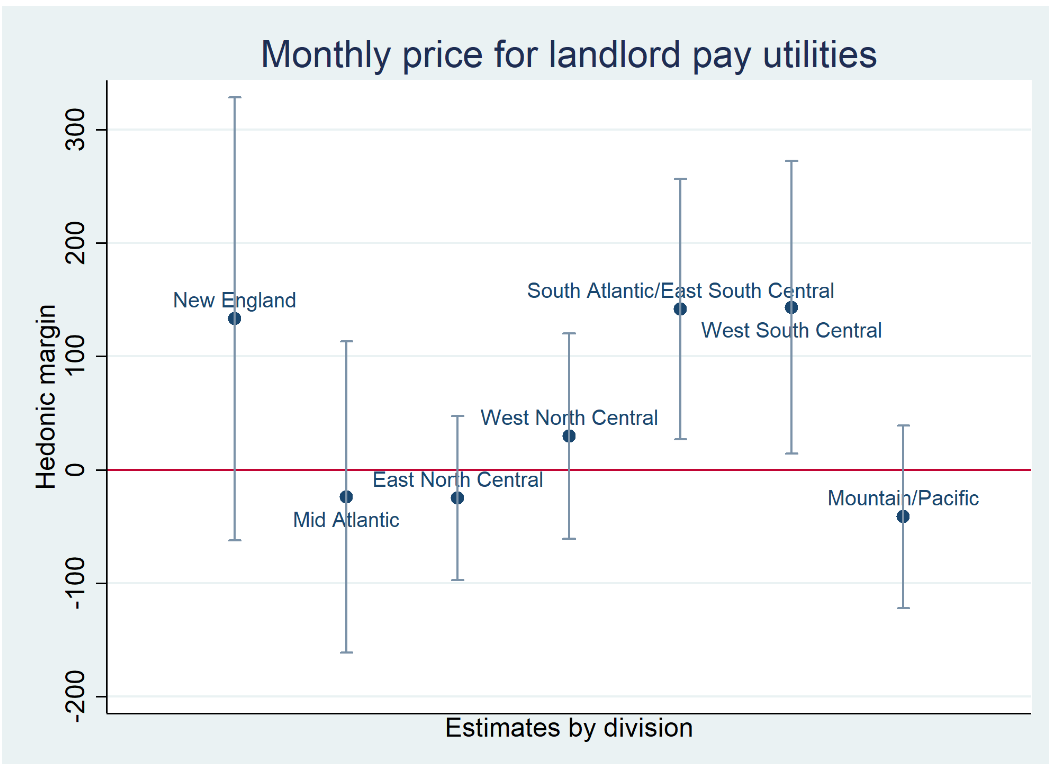

Figure A.4: Plots of estimated average markup for having the landlord-pay heating $\frac{\Delta p}{\Delta R_j}$ by census division from the preferred correlated-random-effects regression in table A.3. 95 percent confidence intervals are derived using cluster-robust standard errors.

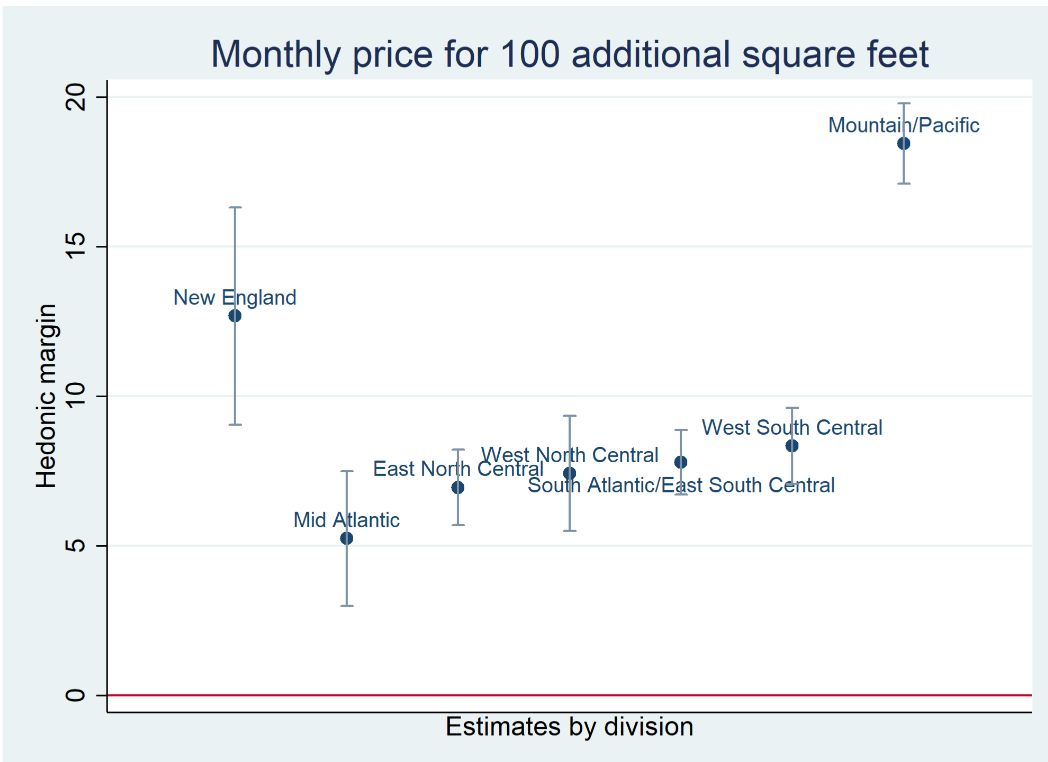

Figure A.5: Plots of estimated average hedonic price for 100 square feet $\frac{\partial \hat{p}}{\partial x_j}$ by census division from the preferred correlated-random-effects regression in table A.3. 95 percent confidence intervals are derived using cluster-robust standard errors.

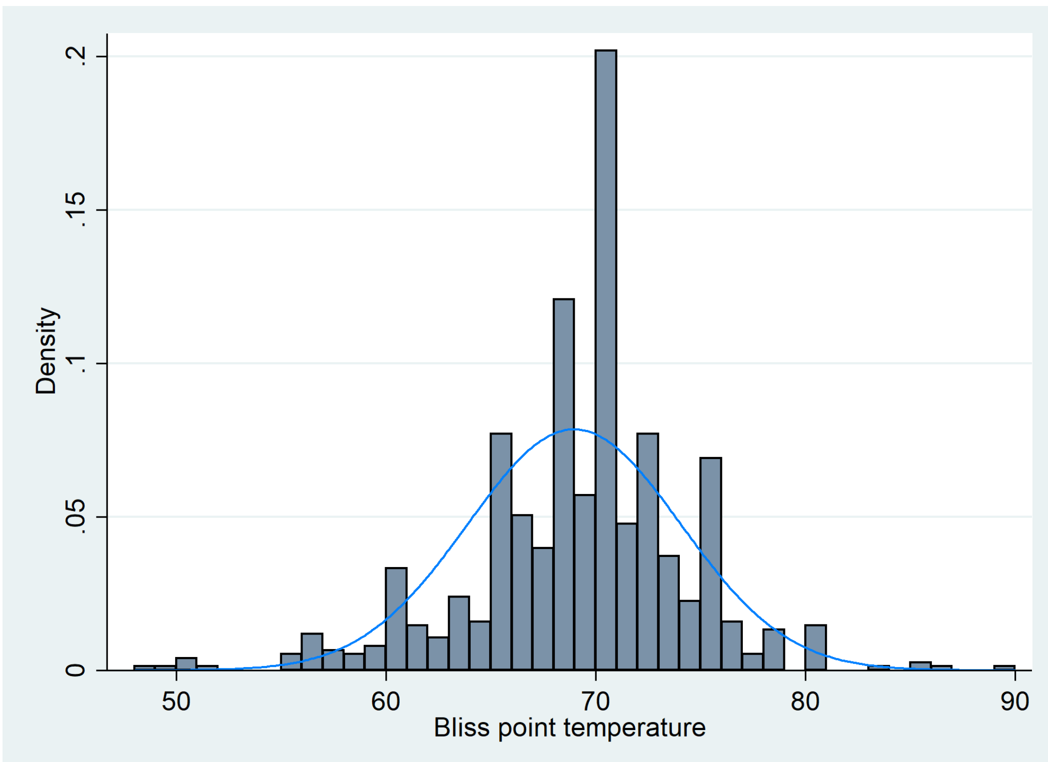

Figure A.6: Density of bliss point temperature settings. Does not include 11 (1.2%) households that did not use heat in the winter. Plotted against a normal distribution.

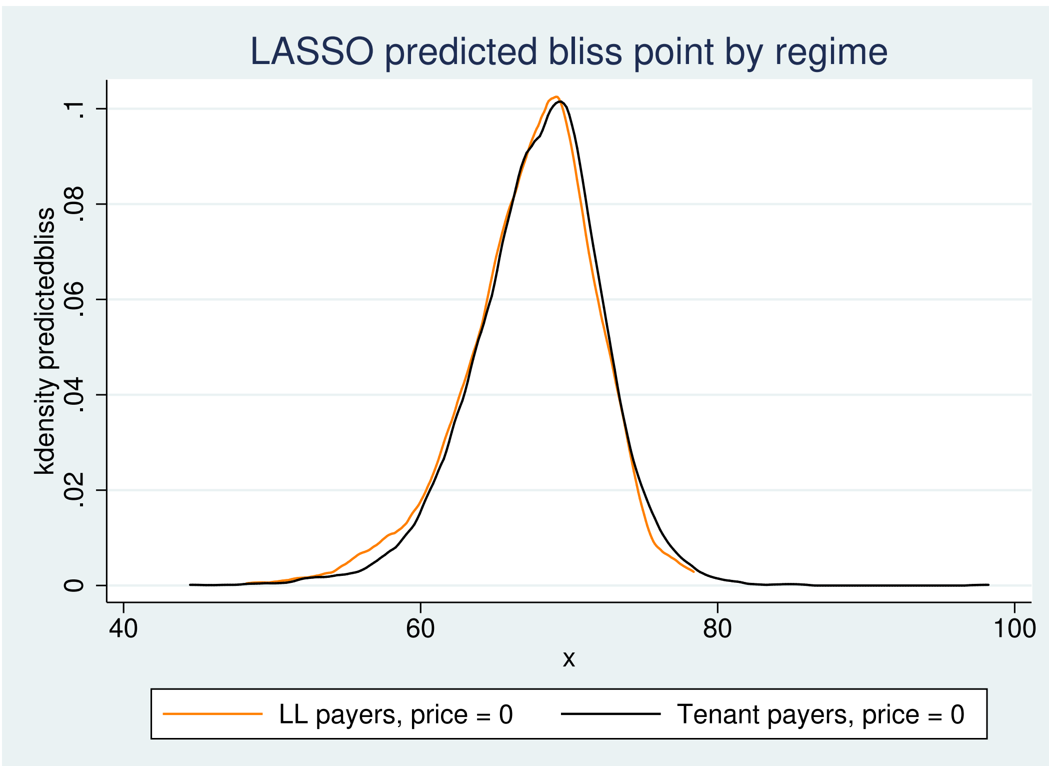

Figure A.7: The LASSO-predicted temperature settings when price is zero (bliss point) for landlord-pay and tenant-pay households. Landlord-pay households are predicted as having lower bliss point temperature preferences.

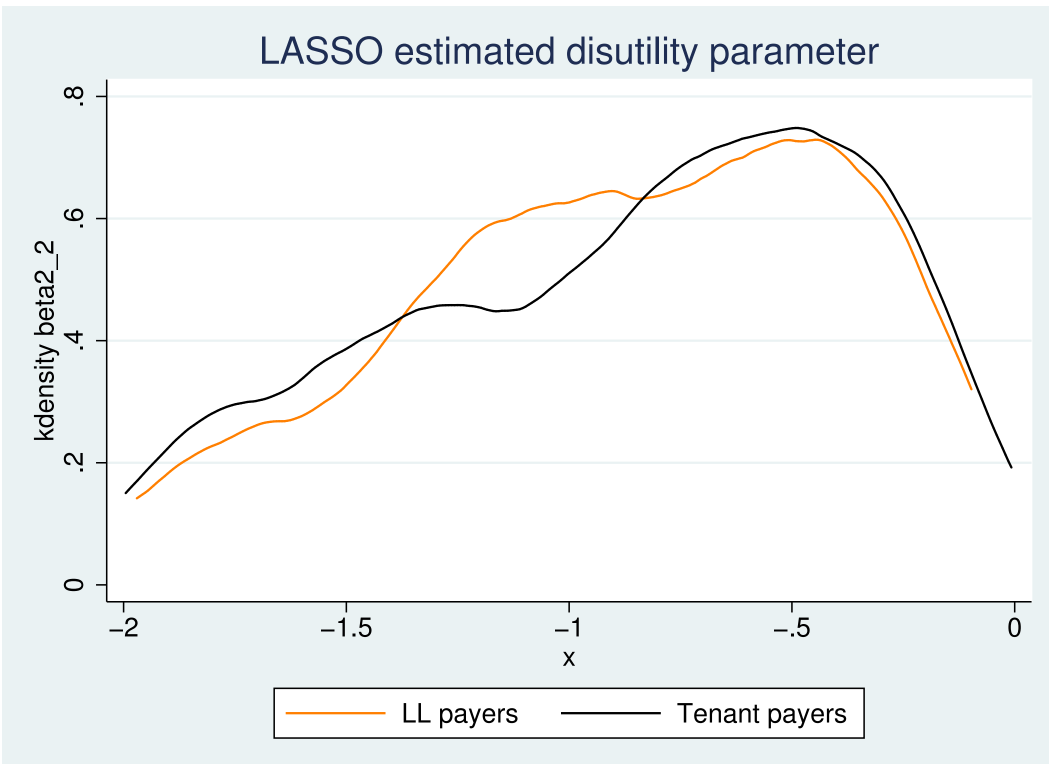

Figure A.8: The implied temperature disutility parameters $\beta_{2,i}$ estimated from the LASSO regression.

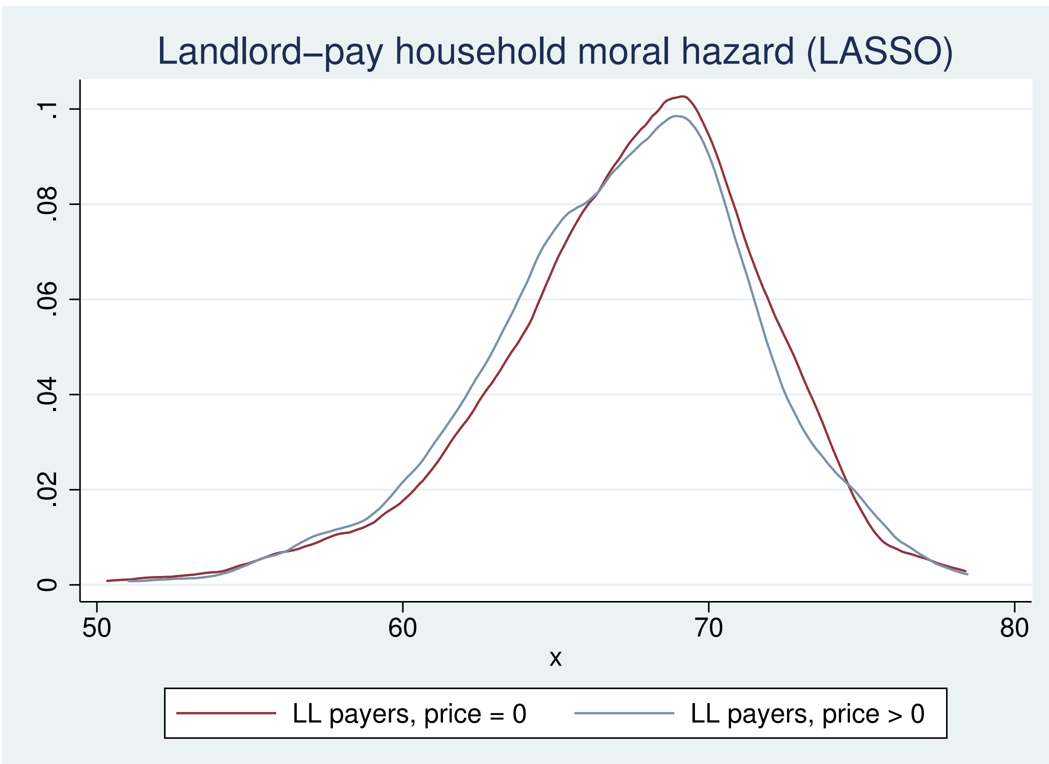

Figure A.9: The LASSO-predicted effect on temperature setting of requiring landlord-pay households pay their own heating bills.

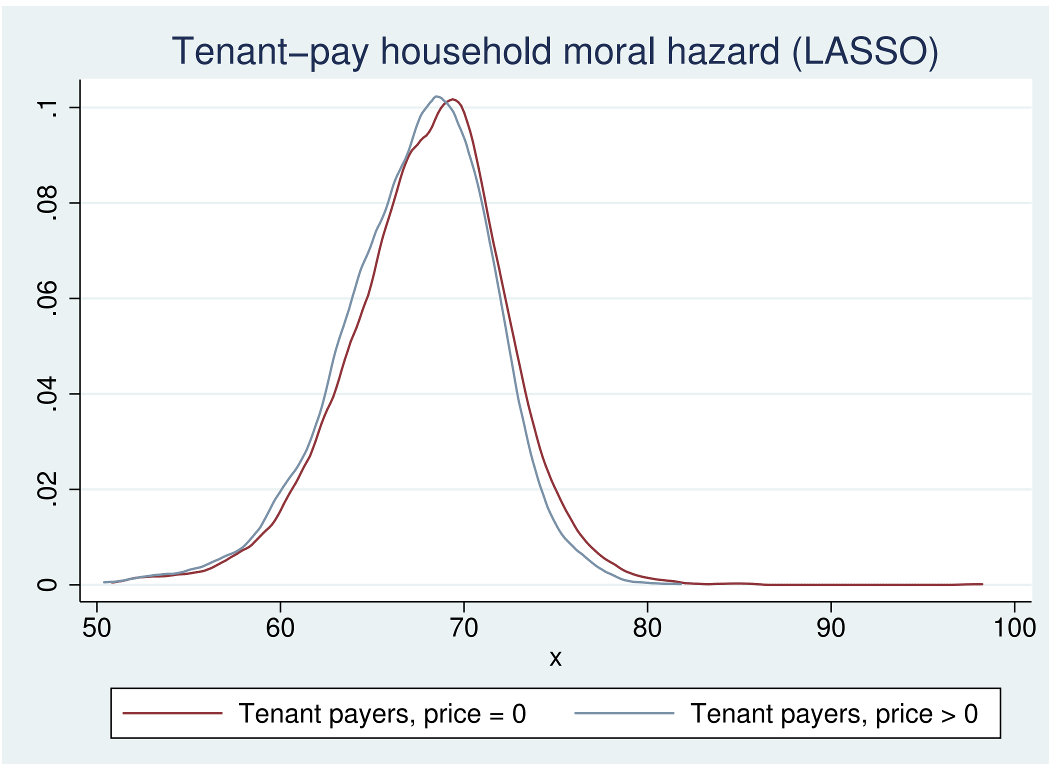

Figure A.10: The LASSO-predicted effect on temperature setting of moving tenant-pay households to a landlord-pay regime.

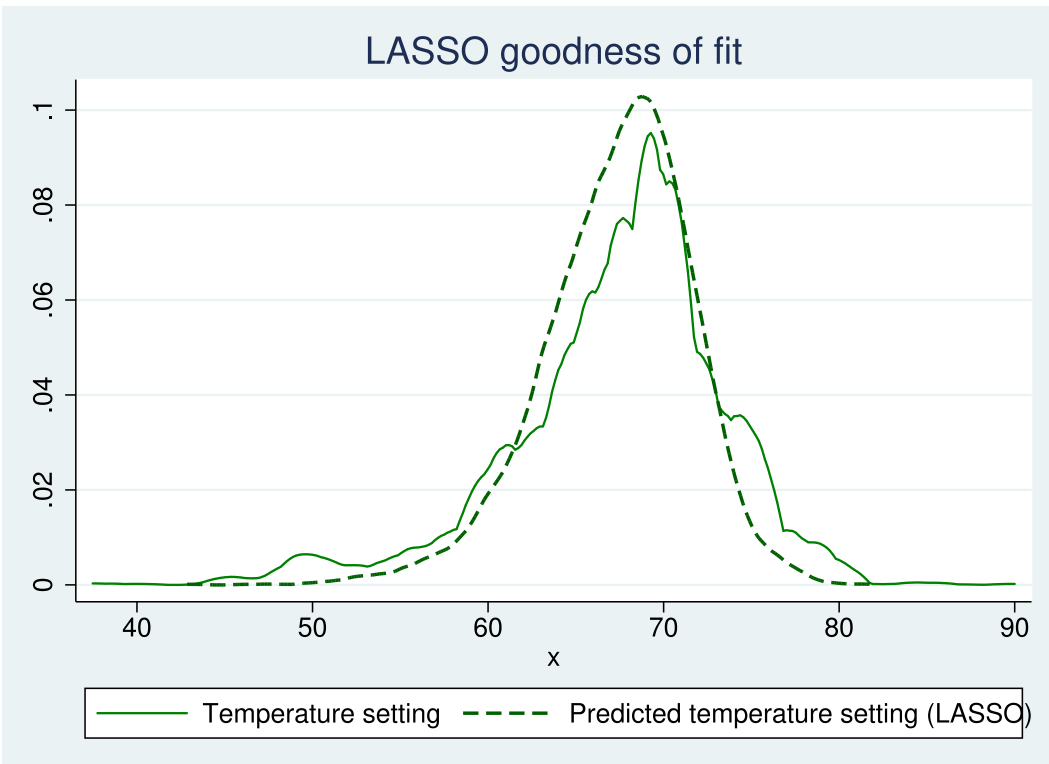

Figure A.11: The predicted temperature setting using LASSO versus the observed temperature setting used to train the algorithm.

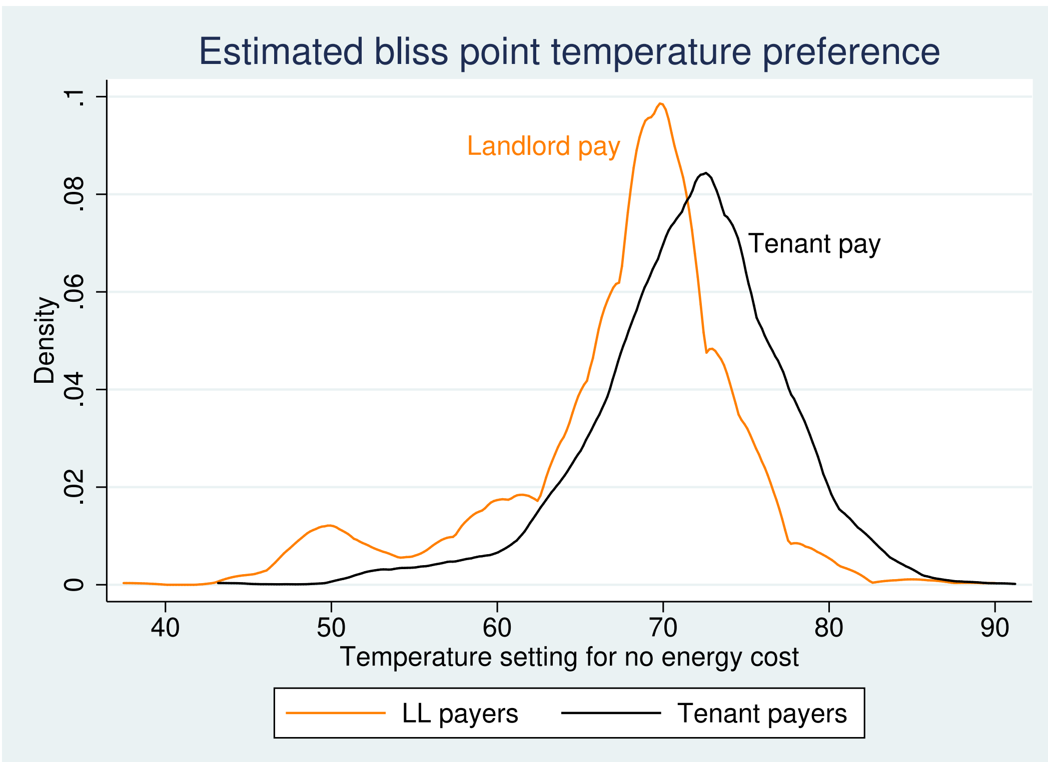

Figure A.12: Estimated bliss points using the preferred $\beta_2 = -0.79$ by regime.

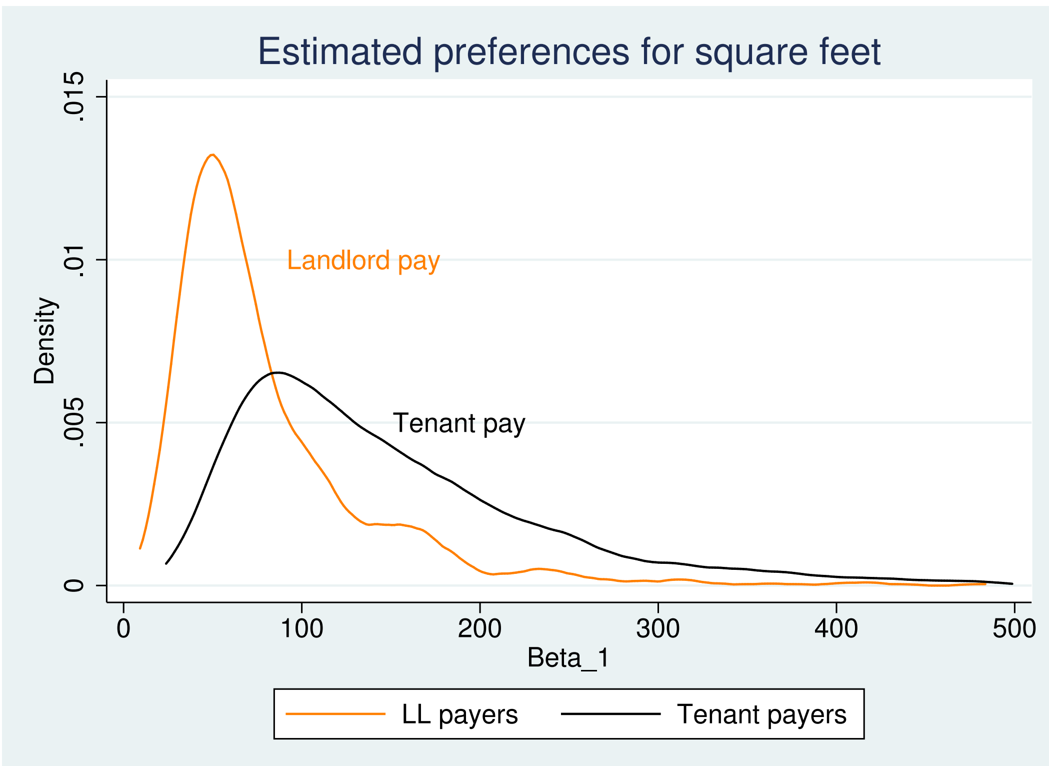

Figure A.13: Kernel density of the distribution of estimated preferences for square feet by regime.

Table A.1: Residential Energy Consumption Survey sample means

	(1) Tenant-pay renters		(2) Landlord-pay renters	
	Mean	Std dev	Mean	Std dev
Mean winter heat temp	67.36	(5.88)	67.40	(7.07)
Temp when home	70.12	(4.91)	70.34	(4.67)
Temp when gone	66.38	(7.02)	68.02	(5.94)
Temp at night	68.59	(5.83)	68.84	(5.69)
Household size	2.65	(1.57)	2.09	(1.38)
Age	40.56	(15.97)	46.59	(20.16)
Female	0.43	(0.49)	0.40	(0.49)
White	0.68	(0.47)	0.70	(0.46)
Black	0.20	(0.40)	0.19	(0.39)
Native American	0.02	(0.13)	0.01	(0.08)
Asian	0.04	(0.20)	0.05	(0.22)
Pacific Islander	0.00	(0.07)	0.00	(0.05)
Other	0.03	(0.17)	0.04	(0.19)
Multi-racial	0.02	(0.15)	0.01	(0.12)
Unemployed	0.29	(0.45)	0.39	(0.49)
Part-time	0.17	(0.38)	0.19	(0.39)
Income \$0 to \$4,999	0.05	(0.22)	0.09	(0.28)
Income \$5,000 to \$9,999	0.08	(0.28)	0.16	(0.37)
Income \$10,000 to \$14,999	0.09	(0.29)	0.16	(0.36)
Income \$15,000 to \$19,999	0.09	(0.28)	0.10	(0.29)
Income \$20,000 to \$29,999	0.17	(0.38)	0.16	(0.36)
Income \$30,000 to \$39,999	0.15	(0.35)	0.11	(0.32)
Income \$40,000 to \$49,999	0.12	(0.32)	0.08	(0.27)
Income \$50,000 to \$74,999	0.14	(0.35)	0.09	(0.29)
Income \$75,000 to \$99,000	0.06	(0.24)	0.03	(0.16)
Income \$100,000 or more	0.05	(0.21)	0.03	(0.17)
New England	0.05	(0.23)	0.14	(0.34)
Middle Atlantic	0.07	(0.26)	0.21	(0.41)
East North Central	0.09	(0.29)	0.17	(0.38)
West North Central	0.10	(0.30)	0.09	(0.28)
South Atlantic	0.16	(0.37)	0.10	(0.30)
East South Central	0.06	(0.23)	0.03	(0.17)
West South Central	0.14	(0.34)	0.10	(0.30)
Mountain	0.09	(0.28)	0.07	(0.25)
Pacific	0.24	(0.43)	0.11	(0.31)
Observations	3,164		829	

Sample means for tenant- and landlord-pay units using electric or gas heating in the Residential Energy Consumption Survey. Standard deviations in parentheses.

Table A.2: American Housing Survey sample means

	(1) Tenant-pay units		(2) Landlord-pay units	
	Mean	Std dev	Mean	Std dev
Monthly rent	803.01	(388.74)	701.65	(401.02)
Energy price per MMBTU	23.07	(10.45)	19.29	(10.18)
Freq rent payment	11.99	(1.62)	12.25	(3.42)
Sq ft	1098.00	(1133.66)	876.79	(1139.78)
# Bedrooms	1.98	(0.89)	1.50	(0.84)
# Baths	1.29	(0.49)	1.09	(0.35)
# Rooms	4.42	(1.32)	3.70	(1.22)
Sq ft	1098.00	(1133.66)	876.79	(1139.78)
Neighborhood rating	7.48	(2.02)	7.45	(2.16)
Unit age (years)	40.76	(23.14)	47.48	(21.73)
# units in building	16.79	(49.95)	48.28	(89.50)
Urban	0.58	(0.49)	0.62	(0.49)
Laundry	0.52	(0.50)	0.19	(0.39)
Central air	0.60	(0.49)	0.43	(0.50)
Furnace heat	0.61	(0.49)	0.53	(0.50)
Gas heat	0.44	(0.50)	0.69	(0.46)
Unit too cold last year	0.09	(0.28)	0.09	(0.29)
Observations	88,700		17,371	

Sample means for tenant- and landlord-pay units using electric or gas heating in the American Housing Survey. Standard errors in parentheses.

Table A.3: Hedonic price estimates using AHS data

$Y = rent_j$	(1) OLS	(2) RE	(3) FE	(4) CRE
Price, 100 sq ft				
New England	13.15	13.87		12.69
	[10.86,15.43]	[10.12,17.61]		[9.044,16.33]
Mid Atlantic	6.874	5.423		5.254
	[5.466,8.282]	[3.208,7.638]		[3.025,7.483]
East North Central	7.416	7.960		6.957
	[6.283,8.549]	[6.649,9.271]		[5.690,8.223]
West North Central	8.223	8.030		7.426
	[6.457,9.989]	[6.039,10.02]		[5.499,9.353]
West South Central	9.148	10.05		8.340
	[8.010,10.29]	[8.720,11.38]		[7.065,9.616]
South Atlantic/	7.888	8.961		7.740
East South Central	[7.049,8.727]	[7.883,10.04]		[6.671,8.809]
Mountain/Pacific	19.43	20.40		18.42
	[18.59,20.27]	[19.10,21.70]		[17.08,19.76]
Landlord-pay markup				
New England	-79.09	14.78	123.5	128.4
	[-192.2,34.05]	[-112.3,141.9]	[-69.39,316.4]	[-66.60,323.4]
Mid Atlantic	-83.20	-54.54	-26.90	-19.69
	[-178.4,12.02]	[-152.2,43.08]	[-160.0,106.1]	[-156.9,117.5]
East North Central	5.812	-11.72	-21.13	-24.39
	[-56.59,68.22]	[-68.37,44.93]	[-92.44,50.18]	[-96.58,47.81]
West North Central	18.12	12.97	31.08	29.03
	[-63.72,99.96]	[-55.75,81.69]	[-58.80,121.0]	[-61.33,119.4]
West South Central	39.63	81.98	148.4	141.1
	[-53.41,132.7]	[-7.716,171.7]	[19.40,277.4]	[12.11,270.1]
South Atlantic/	92.19	115.0	146.9	141.5
East South Central	[24.66,159.7]	[37.85,192.2]	[32.83,260.9]	[27.16,255.8]
Mountain/Pacific	-159.8	-102.6	-37.71	-41.19
	[-209.0,-110.7]	[-162.4,-42.77]	[-118.2,42.79]	[-121.7,39.33]
Unit characteristics	Y	Y	Y	Y
Vintage × year	Y	Y		Y
Unit FE/means			Y	Y
LL pay × year	Y	Y	Y	Y
Observations	107,725	107,725	107,725	107,725

95% confidence intervals in brackets constructed with cluster-robust standard errors. Interpretation is the additional monthly rent for 100 additional square feet and the additional monthly rent to have landlord-pay utilities (i.e., $\hat{\alpha}_{3,d}\bar{P}_{e,d}$), ceteris paribus. Unit characteristics controlled for are bedrooms, bathrooms, air conditioning type, in-unit laundry, number of units in the building, frequency of the rent payment, heating equipment type (e.g., furnace, electric room heaters, etc), non-heating utilities that the landlord pays, whether the unit was too cold last year, a subjective rating of neighborhood quality, an indicator for gas heat, MSA, and division interacted with fuel type.

Table A.4: Fuel use estimates using RECS data

	Natural gas		Electric	
	(1)	(2)	(3)	(4)
$Y = ln(Fuelquantity)$	OLS	2SLS	OLS	2SLS
IHS($Temp_{indoor} - Temp_{outdoor}$)	0.100	0.49	0.019	0.20
	[0.071,0.13]	[0.39,0.59]	[-0.0029,0.042]	[0.079,0.33]
ln(Square feet)	0.21	0.21	0.28	0.26
	[0.14,0.28]	[0.14,0.29]	[0.21,0.36]	[0.14,0.39]
ln(Units in building)	-0.14	-0.15	-0.057	-0.038
	[-0.18,-0.10]	[-0.19,-0.10]	[-0.081,-0.033]	[-0.078,0.0015]
ln(Floors in building)	0.042	0.027	-0.0035	-0.021
	[0.0041,0.080]	[-0.018,0.072]	[-0.028,0.021]	[-0.056,0.015]
Observations	1653	1653	1511	1511
Mean heat bill ($/month)	58.8	58.8	91.0	91.0
Unit characteristics	Yes	Yes	Yes	Yes
Appliances	Yes	Yes	Yes	Yes
VintageXYear	Yes	Yes	Yes	Yes

Interpretation: For a 1% increase in x, there is a $\hat{b}$% increase in monthly heating bill for a fixed energy price, ceteris paribus. Fuel prices at the division level and yearly heating and cooling degree days are instruments. IHS is the inverse hyperbolic sine function, which approximates the natural log, but is defined for non-positive values. 95% Confidence intervals calculated using heteroskedasticity-robust standard errors.

Table A.5: Average marginal effects on monthly heating costs using RECS data

	(1) Natural gas	(2) Electric
Heat setting - one degree change	3.57	2.89
	[2.76,4.37]	[1.06,4.71]
Square feet - 100 sq ft change	1.01	2.42
	[0.63,1.39]	[1.31,3.53]
Units in building - one unit change	-5.86	-1.74
	[-7.71,-4.01]	[-3.61,0.13]
Floors in building - one floor change	1.11	-1.18
	[-0.80,3.02]	[-3.25,0.90]
Observations	1653	1511
Unit characteristics	Yes	Yes
Appliances	Yes	Yes
VintageXYear	Yes	Yes

Interpretations: For a unit change in x, there is a $\hat{b}$ dollar increase in monthly heating bill, ceteris paribus. Estimates from the two-stage-least-squares specification in table A.4. 95% confidence intervals are bootstrapped using 1,000 replications.

Table A.6: Temperature setting estimation using RECS data

	OLS				2SLS			
$Y = temp\ setting$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\hat{\beta}_2$	-0.56	-0.57	-0.56	-0.53	-0.98	-1.06	-0.36	1.44
	[-0.73,-0.38]	[-0.75,-0.39]	[-0.73,-0.39]	[-0.69,-0.37]	[-9.43,7.46]	[-13.0,10.9]	[-2.86,2.14]	[-8.14,11.0]
$\hat{\delta}$	-1.76	-1.65	-1.29	-1.15	-3.14	-3.16	-0.69	4.92
	[-2.36,-1.16]	[-2.24,-1.06]	[-1.89,-0.70]	[-1.73,-0.58]	[-32.1,25.8]	[-42.0,35.7]	[-8.94,7.56]	[-24.6,34.4]
$\hat{s}_i^b \| R_{j*} = 0$	67.4	67.4	67.4	67.4	67.4	67.4	67.4	67.4
	[66.9,67.9]	[66.9,67.9]	[66.9,67.9]	[66.9,67.9]	[66.9,67.9]	[66.9,67.9]	[66.9,67.9]	[66.9,67.9]
$\hat{s}_i^b \| R_{j*} = 1$	69.2	69.2	69.2	69.1	70.5	70.8	68.5	62.7
	[68.7,69.6]	[68.8,69.6]	[68.7,69.6]	[68.6,69.5]	[41.7,99.4]	[30.2,111.4]	[59.3,77.8]	[31.6,93.8]
Observations	3993	3993	3993	3993	3993	3993	3993	3993
First stage F					29.6	21.0	35.3	4.96
HH characteristics		Yes	Yes	Yes		Yes	Yes	Yes
Division			Yes	Yes			Yes	Yes
Year				Yes				Yes

Estimation of equation 1.18. $\hat{\beta}_2$ is an estimate of the mean disutility parameter for deviation from the bliss point. Estimated using a two-step estimator for marginal cost. Instrument for marginal cost in 2SLS estimates is division fuel price. 95% confidence intervals including the sampling error for marginal cost bootstrapped with 1,000 replications. Sample of all renters from the RECS.

Table A.7: Temperature setting estimation using choice experiment

	(1)	(2)
	OLS	FE
$\hat{\beta}_2$	-0.76	-0.79
	[-1.08,-0.43]	[-0.97,-0.61]
Constant	-0.45	-0.42
	[-0.77,-0.14]	[-0.60,-0.24]
Observations	1242	1242
FE		Yes

Estimation of equation 1.23. 95 percent confidence intervals in brackets. Estimated on a sample of choice experiment respondents for which hypothetical marginal cost is varied randomly. I conducted the choice experiment in March 2018 using the Qualtrics Panel and a nationally representative sample of US individuals.

Table A.8: LASSO heterogeneity

	OLS	LASSO
Mean predicted bliss temp	72.01	67.82
Std dev.	5.56	4.39
Mean predicted tenant pay temp	67.31	67.31
Std dev.	4.44	4.24
Mean implied β_2	-1.54	-9.31
Std dev.	10.57	59.29
RMSPE	8.73	5.71
MSPE	76.15	32.55
Bias	-0.01	0.00
Covariates	1590	1590
FinalCovariates	1590	1209
Observations	3993	3993

Mean-squared error, root-mean-squared error, and bias statistics calculated using ten-fold cross validation.

Table A.9: Simulation results:
Effect of changing all landlord-pay contracts to tenant-pay

	(1)	(2)	(3)
Reduction in energy expenditures per household, per year			
Temperature reduction	4.15	6.54	3.27
Square feet reduction	139.57	139.57	139.57
Implied welfare loss from landlord pay per household, per year			
Moral hazard DWL	\$7.83	\$12.33	\$6.17
Sorting DWL	\$1.34	\$1.34	\$1.34
Carbon damages	\$3.64	\$5.46	\$2.97
Total US, per year average			
US yearly welfare loss	\$839,000,000	\$1,260,000,000	\$682,000,000
β_2	-0.79	-0.5	-1.00

Average household responses to eliminating moral hazard by requiring tenant-pay heating in all units. Total welfare losses include deadweight losses from moral hazard, equilibrium sorting, and external cost of carbon (\$40/ton CO2) over all landlord-pay units using survey weights.

APPENDIX B

SUPPLEMENTAL INFORMATION FOR CHAPTER 1

B.1 Science and engineering findings on temperature preferences

The science and engineering literatures characterize temperature preferences along three dimensions: temperature sensitivity, temperature discomfort, and mean temperature preference (Kingma and van Marken Lichtenbelt, 2015; Karjalainen, 2012; Indraganti and Rao, 2010; Schellen et al., 2010). Temperature sensitivity is the ability to *detect* changes in temperature from a baseline. Temperature discomfort is a subjective rating of how uncomfortable deviations from a bliss point temperature are. Finally, mean temperature preference refers to the temperature bliss point itself. In terms of the model in this paper, temperature sensitivity and discomfort are factors in the disutility parameter (β_2), while mean temperature preference is a measure of bliss point s_i^b (where the indoor temperature's contribution to utility is $(-1/\beta_2)(s_i - s_i^b)^2$).

Human temperature preferences depend on several physiological factors. Diets rich in Vitamin C improve extreme heat tolerance in many studies (Ringsdorrf Jr. and Cheraskin, 1982). Older individuals prefer higher temperatures than younger individuals (Schellen et al., 2010), though the elderly find colder temperatures less unpleasant perhaps due to decreased sensitivity (Taylor et al., 1995). The majority of scientific temperature studies conclude that women prefer higher temperatures and are more sensitive to temperature extremes (Kingma and van Marken Lichtenbelt, 2015; Karjalainen, 2012).[1]

Culture may play a role in driving temperature preference differences between sexes. A study of Finnish men and women found that while women report more discomfort than men in both low and high temperatures, women are less likely to adjust the thermostat and are more likely to prefer warmer settings in general (Karjalainen, 2007). In another study, Indian men and women exhibited

[1]See also Karjalainen (2007); Fanger (1970); Parsons (2002); Cena and de Dear (2001); Muzi et al. (1998); Pellerin and Candas (2003); Griefahn and Künemund (2001); Nakano et al. (2002); Nagashima et al. (2002)

a slightly different effect: while women were more sensitive to temperature changes than men, they were less likely to report discomfort (Indraganti and Rao, 2010). This cross-country difference is suggestive of a cultural role in temperature preference. In the United States, women express greater temperature discomfort than men (Beshir and Ramsey, 1981). Another study examining more general aspects of workplace environment such as air quality, social conditions, and noise levels in addition to temperature finds that women report a greater number of work-related health symptoms than men in general (Reynolds et al., 2001). Surprisingly, temperature was a statistically significant determinant of workplace health problems for men and *not* women. The authors speculate that these findings are likely driven by cultural attitudes of “stoic" behavior among men. In any case, sex and culture are both important determinants of heat preference.

When exposed to heat and cold over time, the human body adapts physiologically to become more tolerant of hot and cold temperatures (Young, 2010). Households living in colder climates will likely be more tolerant of cool indoor temperatures. Through this channel, temperature preference may also be related to income if low-income households develop a resilience to cold weather when conserving on heating costs, or if low-income households tend to work or commute outdoors in cold weather. Indraganti and Rao (2010) demonstrate an income-temperature effect for heat resilience in India with less wealthy individuals reporting a higher comfort level in extreme summer weather conditions. Thus, local climate and income will likely influence temperature preference.[2]

Households reveal bliss-point temperature preferences when heating without price constraints. The science and engineering literature has found that physiological factors such as age, sex, and previous temperature exposure influence temperature preferences. Other factors such as culture, behavioral or psychological adaptation (de Dear and Brager, 1998), and idiosyncratic preferences

[2]In theory, this acclimatization process could be used as a “defensive" behavior to mitigate energy expenditures. There is recent experimental evidence that suggests individuals can consciously alter the body’s internal temperature response with training (see Kox et al.’s (2014) study of the “Wim Hof method" in the *Proceedings of the National Academy of Sciences*). This relatively unknown method involves bathing in temperatures near 32°F over weeks to become acclimatized to low temperatures. While this form of defensive behavior might be used in theory to reduce heating costs, I do not think it is likely.

probably play a large role. If temperature preference is largely determined physiologically, then preferences can be predicted using reported demographics and revealed heating choices.

B.2 Data cleaning notes

Vacation homes, vacant units, timeshares, tents, mobile homes, and hotel rooms are dropped from the sample. This essentially leaves only traditional apartment units and rental homes in the sample. In addition, units where no rent is paid, the rent is reported to be adjusted due to relationship with the owner, the frequency of the rent payment is not reported, and units where the square footage is not reported are dropped from the sample. The American Housing Survey is topcoded in rent price, so units with the top 3% of rents in the sample in the National sample are dropped, and the top 3% of rents in each geographic area in the Metropolitan sample are dropped.

B.3 Alternative explanations

I provide descriptive evidence showing that landlord-pay households are in control of their temperature settings. One alternative explanation for the lack of evidence that households select into landlord pay regimes based on bliss point temperature preference is that thermostat settings by landlord-payers may not reflect the true bliss point. If the landlord controls the thermostat in these regimes, interpreting thermostat settings as revealed bliss point is improper; however, there is only limited evidence to support this hypothesis. In the 2015 version of the Residential Energy Consumption Survey, respondents are asked whether they control their thermostat. 86% of landlord-pay households report that they control the temperature setting in their home. The difference in temperature settings between those in control and not in control is not statistically different from zero. Given the small number of households and the lack of perceivable differences in temperature settings, I conclude that it is not likely that this effect is driving the results.

APPENDIX C

FIGURES AND TABLES FOR CHAPTER 2

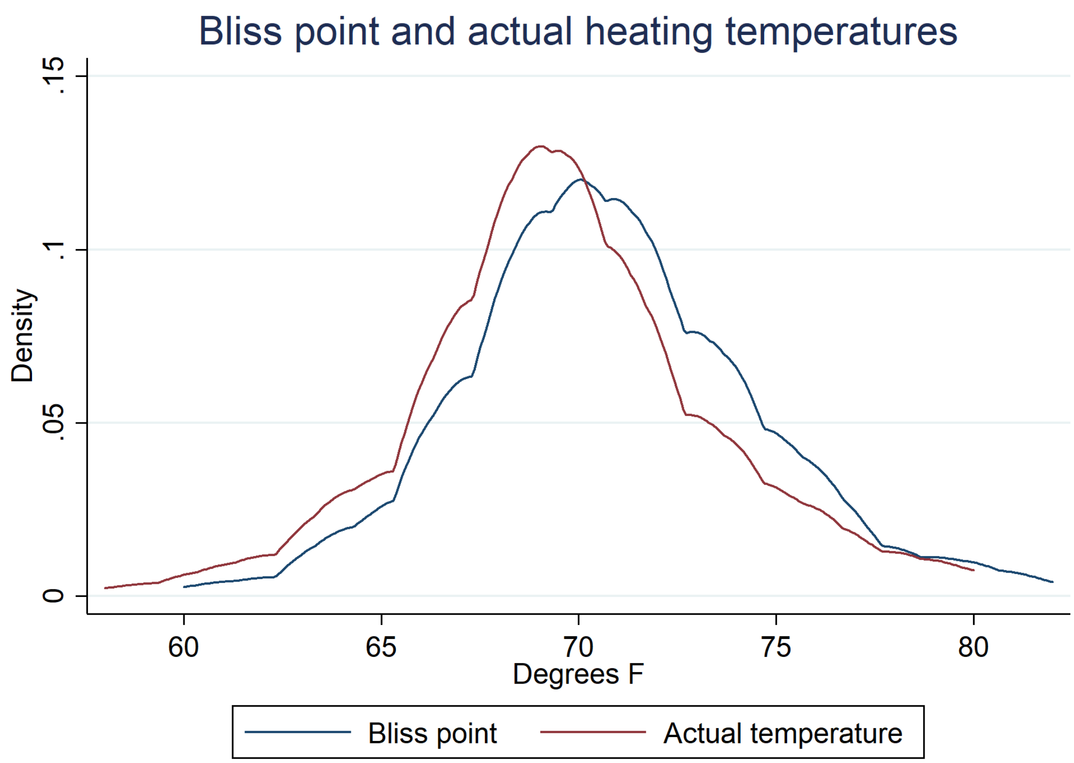

Figure C.1: Kernel density of participants' bliss point temperature preferences and real temperature settings. A Kolmogorov-Smirnov test of equivalence of distributions easily rejects the null hypothesis that the distributions are equal.

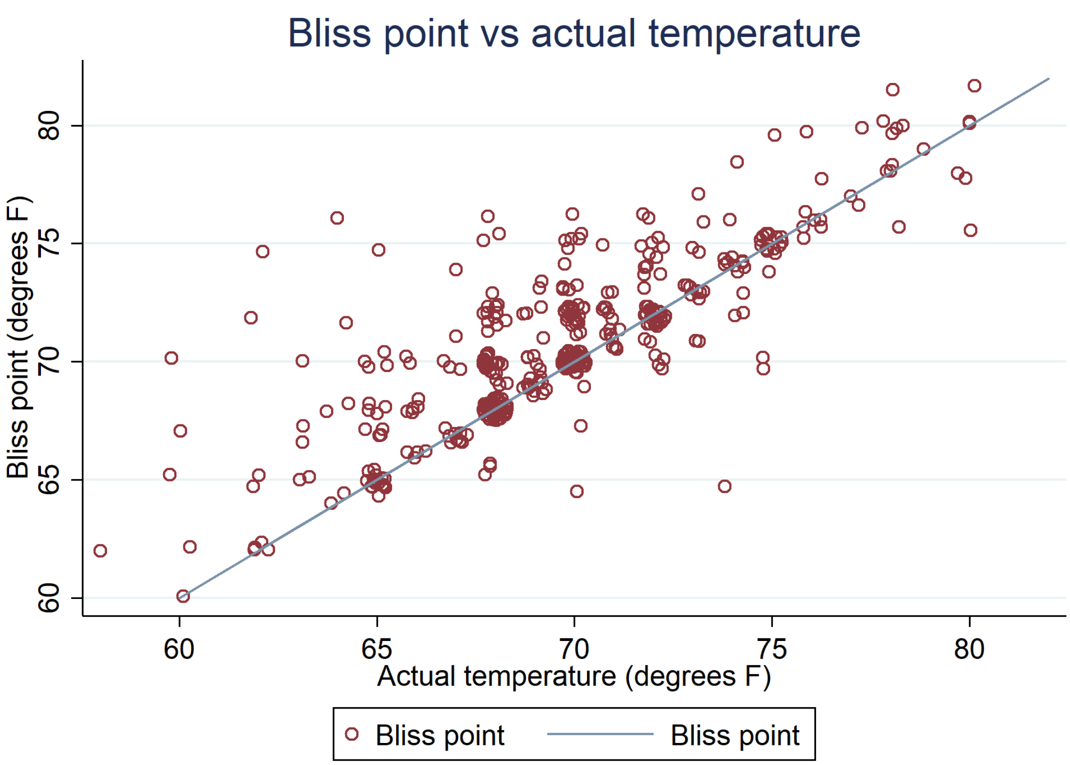

Figure C.2: Scatterplot of bliss point temperature preferences and actual temperature settings with a 45 degree line for reference. 54 percent of respondents set the thermostat equal to the bliss point. 6.5 percent reported setting the thermostat greater than the bliss point. Random noise has been added to the data to show clustering on common temperature choices such as 70 degrees Fahrenheit.

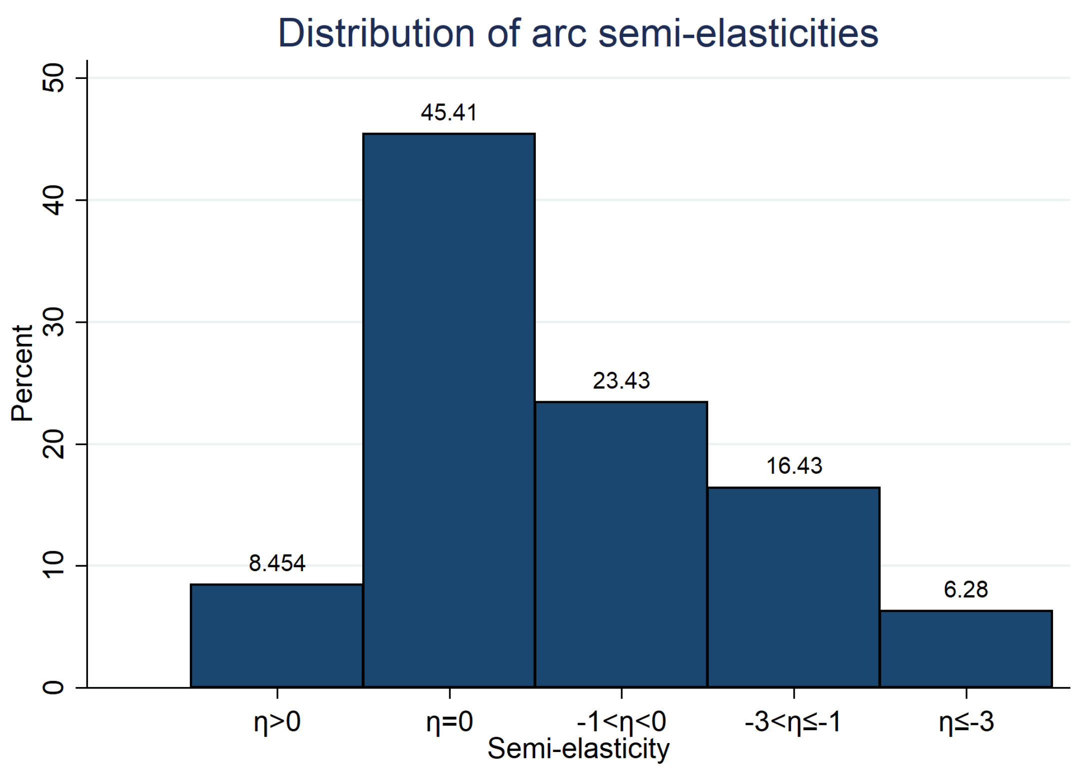

Figure C.3: Histogram of estimated arc semi-elasticities (η).

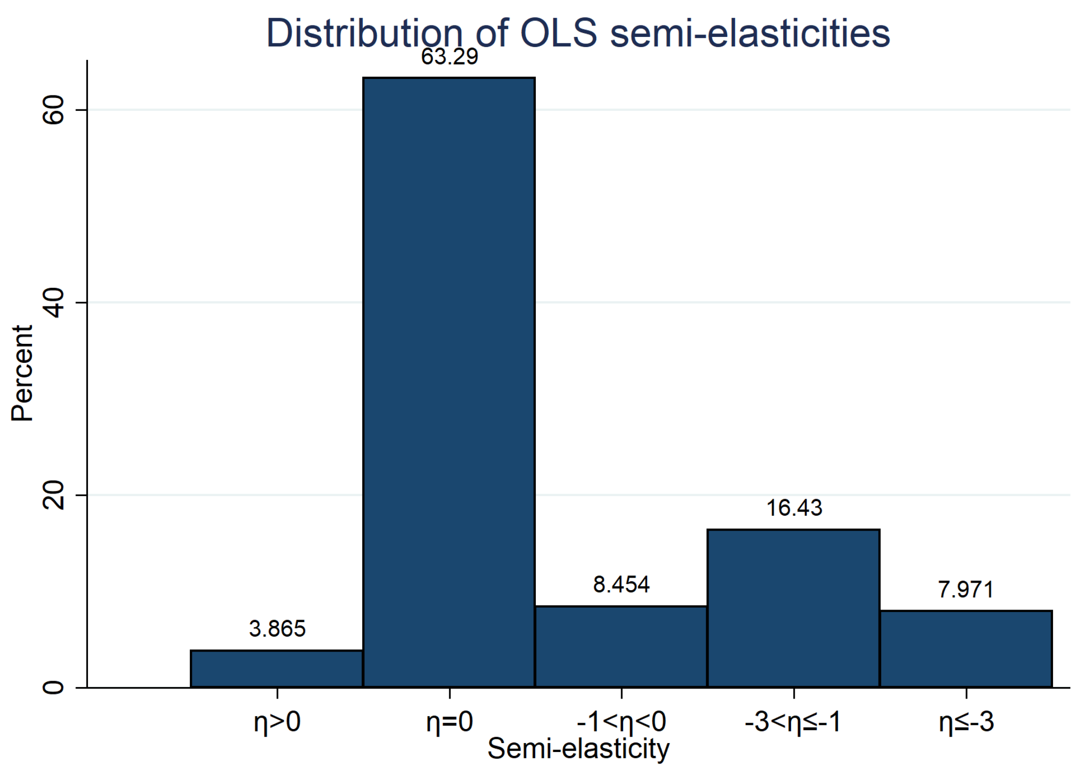

Figure C.4: Histogram of semi-elasticities (η) estimated using OLS.

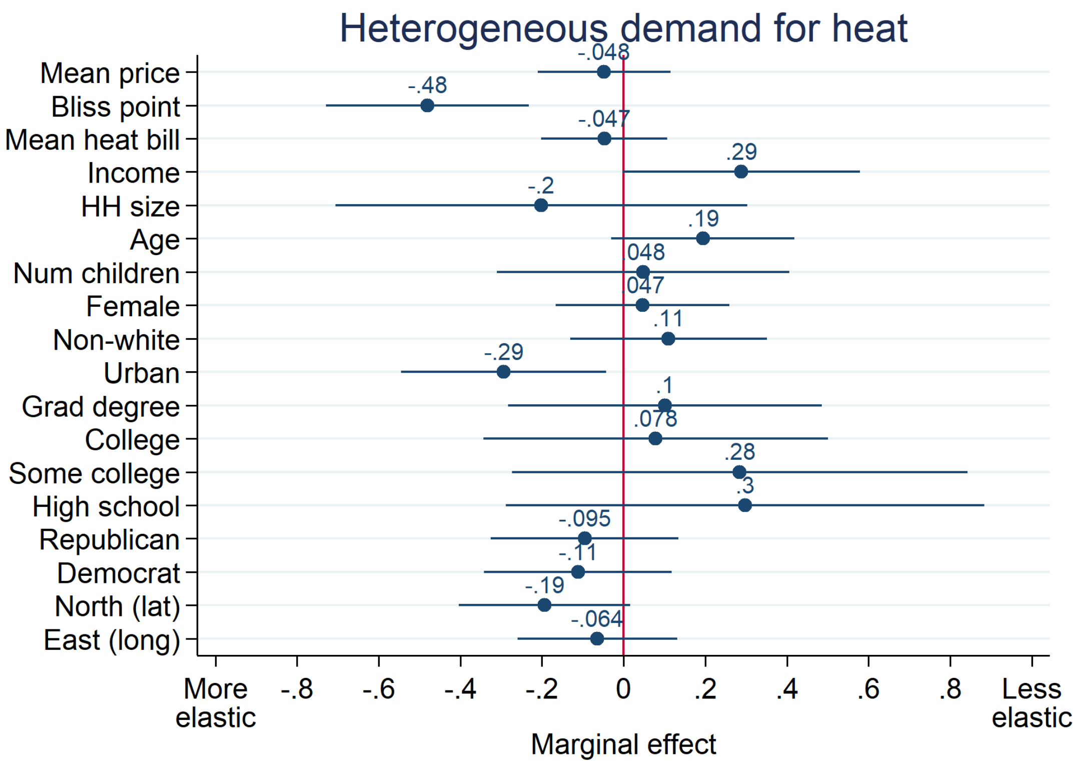

Figure C.5: The marginal effects from a Tobit estimation of the estimated arc semielasticities on average price and participant demographics. The marginal effects from this estimation can be interpreted as the change in arc semi-elasticity for a one-standard-deviation change in the predictor variable holding the other predictor variables constant. 95 percent confidence intervals are bootstrapped using 1000 replications with sampling at the participant level.

Table C.1: Participant sample mean and standard deviation

	(1) Mean	(2) Std dev
Income	79925.13	(71251.68)
Monthly heat bill	121.29	(145.81)
Bliss point	70.77	(3.59)
Temperature at home	69.70	(3.71)
Household size	2.91	(2.85)
Age	50.22	(16.60)
Children	0.44	(1.11)
Female	0.60	(.49)
Non-white	0.29	(.45)
Urban	0.76	(.43)
High school	0.37	(.48)
Some college	0.28	(.45)
College	0.19	(.39)
Graduate degree	0.10	(.30)
Republican	0.28	(.45)
Democrat	0.33	(.47)
Respondents	414	
Observations	1242	

Sample means presented with sample standard deviations in parentheses.

Table C.2: Estimated semi-elasticities.

	(1) Elasticity	(2) Lower bound	(3) Upper bound	(4) N
OLS	-.52	-.73	-.31	1242
FE	-.56	-.69	-.43	1242
Arc semi-elasticity	-.69	-.97	-.43	1242
Individual OLS	-.66	-.86	-.46	1242

The interpretation is for a 100 percent increase in price, the average participant will reduce the thermostat by η degrees. 95 percent confidence intervals bootstrapped using 1000 replications with sampling at the participant level.

APPENDIX D

FIGURES AND TABLES FOR CHAPTER 3

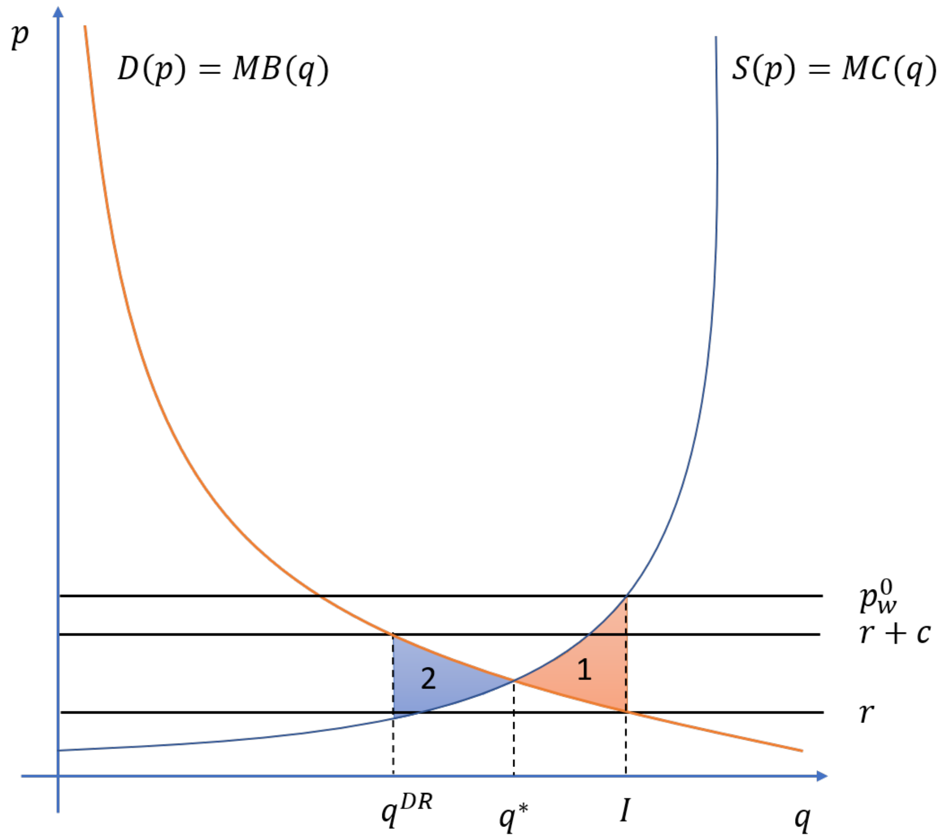

Figure D.1: Aggregate supply and demand curves for electricity consumption. Area 1 corresponds to the surplus loss in the no-demand-response baseline and area 2 corresponds to the surplus loss from a demand response compensation level greater than $p_w - r$.

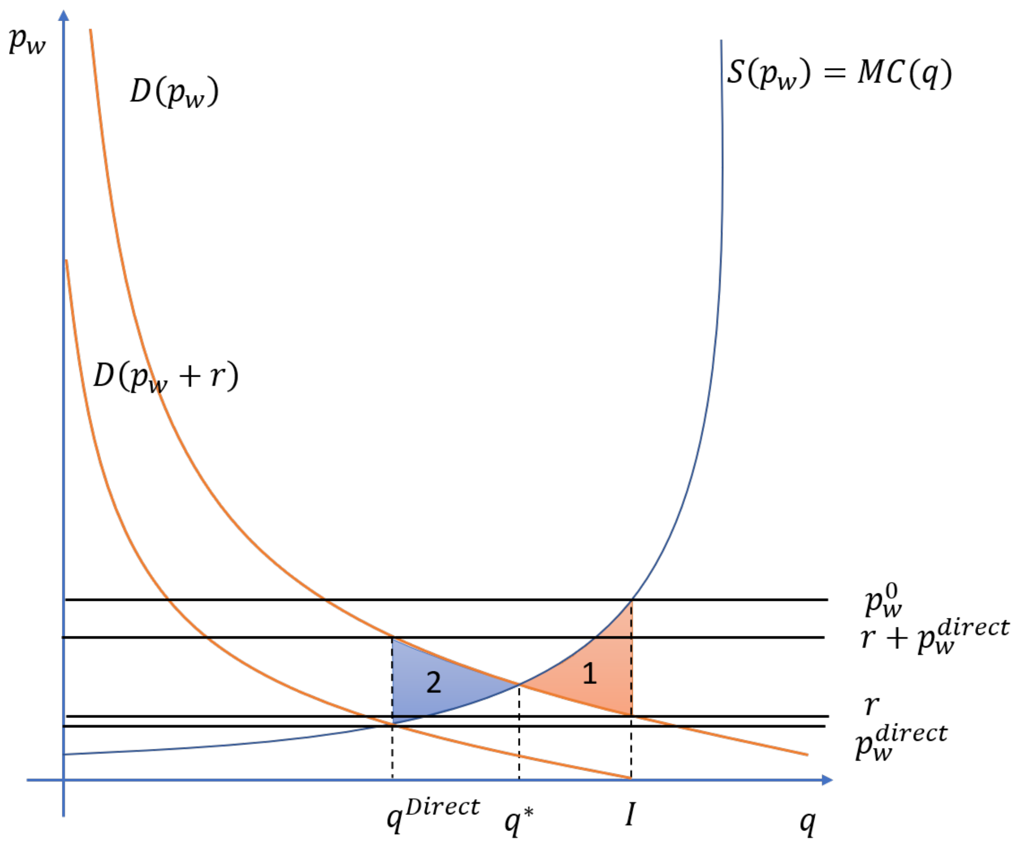

Figure D.2: The equilibrium outcome and surplus loss under direct participation. Area 1 corresponds to the surplus loss under the no-demand-response baseline. When electricity consumers particpate directly in demand response, the opportunity cost already includes the retail rate, so the demand curve shifts downward to $D(p_w + r)$. The intersection of $D(p_w + r) = S(p_w)$ determines the market-clearing wholesale price under direct participation, which decreases. Area 2 corresponds to the surplus loss under direct participation.

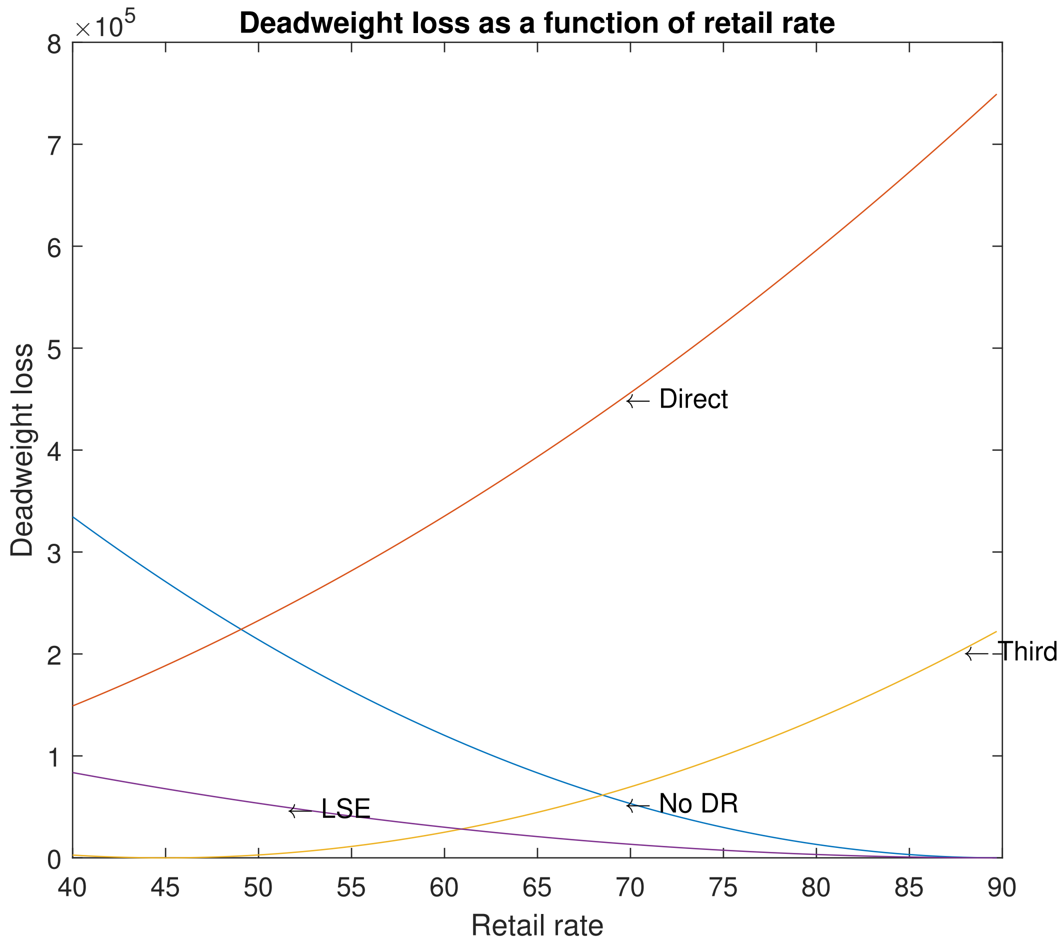

Figure D.3: Simulated deadweight loss by regime type plotted against the demand elasticity parameter.

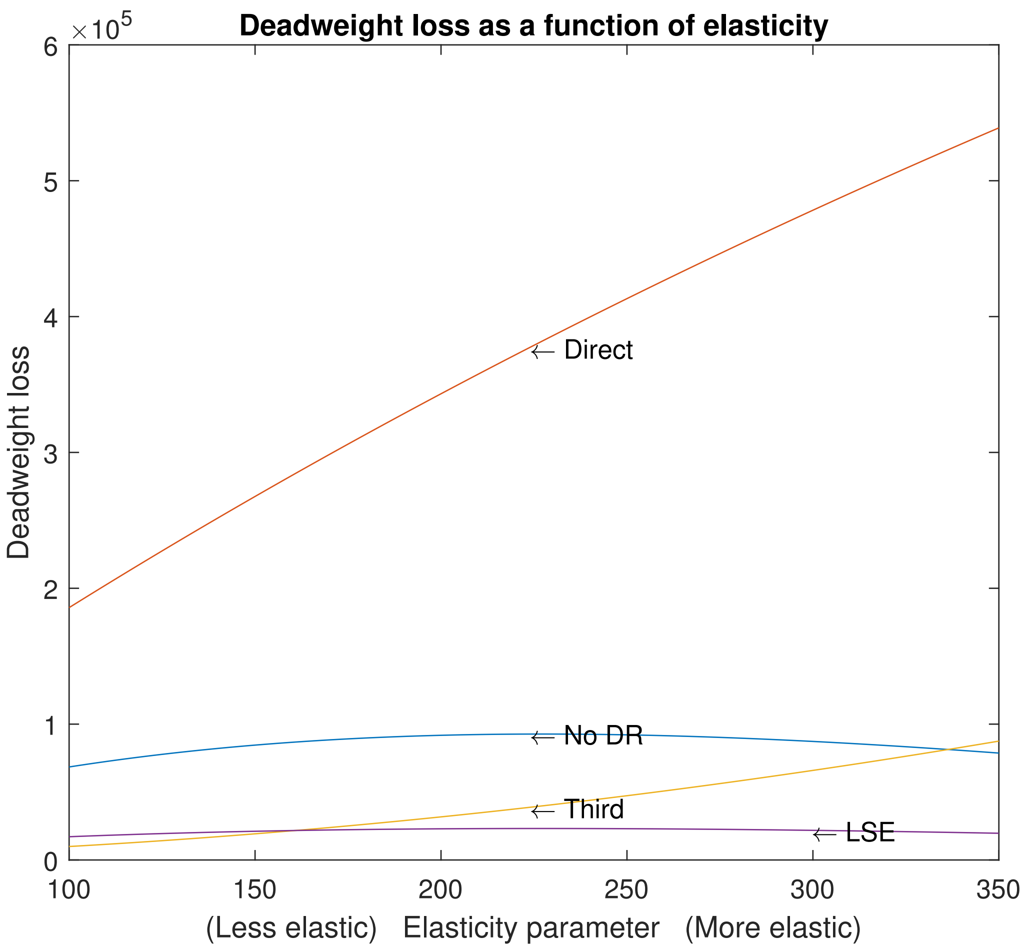

Figure D.4: Simulated deadweight loss by regime type plotted against the retail rate.

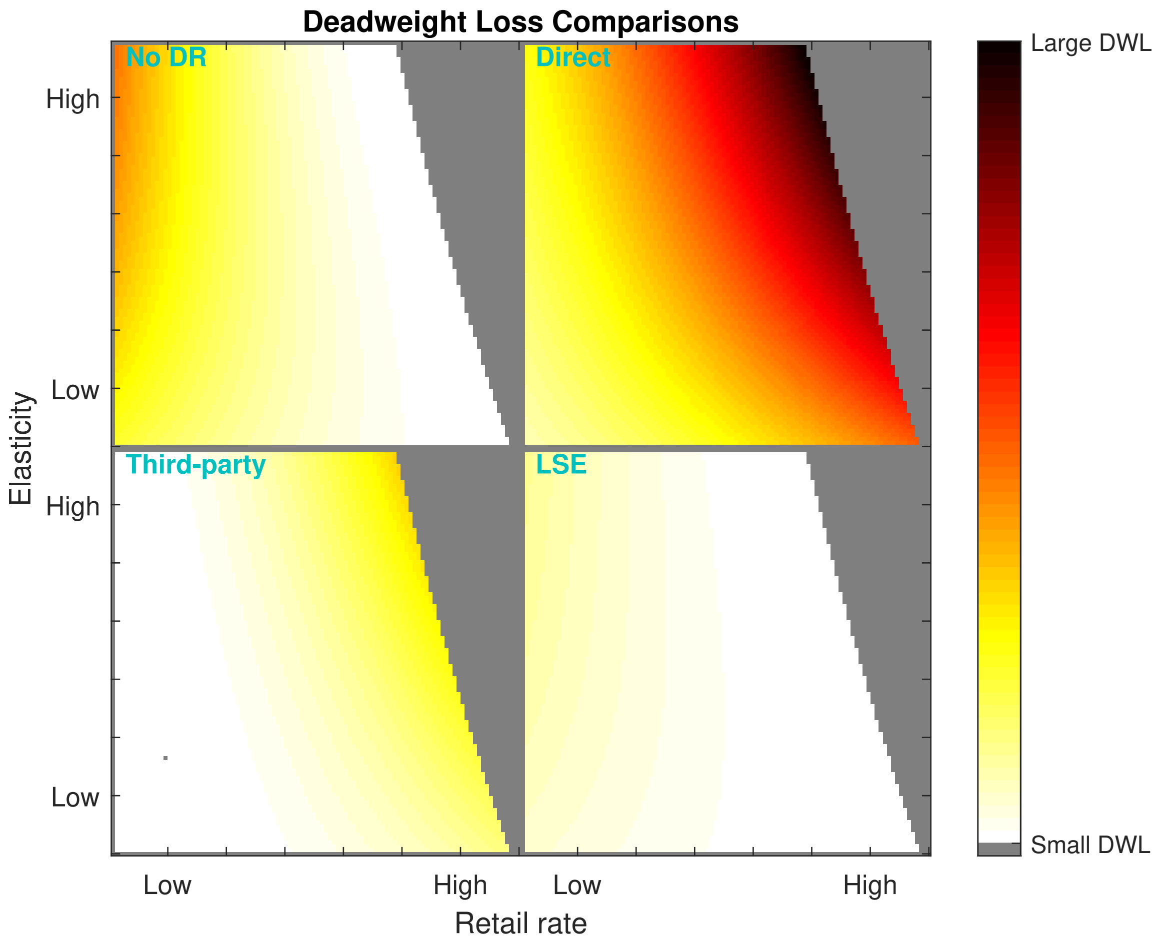

Figure D.5: Simulated deadweight loss for each regime type versus the retail rate and the elasticity parameter. Darker colors indicate larger simulated deadweight losses. The gray regions indicate where the wholesale price is lower than the retail rate.

BIBLIOGRAPHY

BIBLIOGRAPHY

Aggarwal, R. M. (2007). Role of risk sharing and transaction costs in contract choice: Theory and evidence from groundwater contracts. *Journal of Economic Behavior and Organization*, 63(3):475 – 496.

Allcott, H. (2015). Site selection bias in program evaluation. *The Quarterly Journal of Economics*, 130(3):1117–1165.

Allcott, H. and Greenstone, M. (2012). Is there an energy efficiency gap? *Journal of Economic Perspectives*, 26(1):3–28.

Allcott, H. and Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10):3003–37.

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

Allen, D. and Lueck, D. (1992). Contract choice in modern agriculture: Cash rent versus cropshare. *The Journal of Law and Economics*, 35(2):397–426.

Allen, D. W. and Lueck, D. (1995). Risk preferences and the economics of contracts. *The American Economic Review*, 85(2):447–451.

Allen, D. W. and Lueck, D. (1999). The role of risk in contract choice. *The Journal of Law, Economics, and Organization*, 15(3):704–736.

Allen, R. G. D. and Lerner, A. P. (1934). The concept of arc elasticity of demand. *The Review of Economic Studies*, 1(3):226–230.

Auffhammer, M. and Rubin, E. (2018). Natural gas price elasticities and optimal cost recovery under consumer heterogeneity: Evidence from 300 million natural gas bills. Technical report, Energy Institute at Haas.

Autor, D., Duggan, M., and Gruber, J. (2014). Moral hazard and claims deterrence in private disability insurance. *American Economic Journal: Applied Economics*, 6(4):110–41.

Baicker, K., Mullainathan, S., and Schwartzstein, J. (2015). Behavioral hazard in health insurance. *The Quarterly Journal of Economics*, 130(4):1623–1667.

Bajari, P. and Benkard, C. L. (2005). Demand estimation with heterogeneous consumers and unobserved product characteristics: A hedonic approach. *Journal of Political Economy*, 113(6):1239–1276.

Bajari, P. and Kahn, M. E. (2005). Estimating housing demand with an application to explaining racial segregation in cities. *Journal of Business and Economic Statistics*, 23(1):20–33.

BEA (2019). Gross domestic product [gdpa]. Technical report, United States Bureau of Economic Analysis.

Belloni, A. and Chernozhukov, V. (2013). Least squares after model selection in high-dimensional sparse models. *Bernoulli*, 19(2):521–547. arXiv: 1001.0188.

Beshir, M. and Ramsey, J. (1981). Comparison between male and female subjective estimates of thermal effects and sensations. *Applied Ergonomics*, 12(1):29 – 33.

Borck, R. and Brueckner, J. K. (2017). Optimal energy taxation in cities. *Journal of the Association of Environmental and Resource Economists*, 5(2):481–516.

Bowring, J. (2018). 2017 State of the market report for PJM. Technical report, Monitoring Analytics, LLC presentation.

Brot-Goldberg, Z. C., Chandra, A., Handel, B. R., and Kolstad, J. T. (2017). What does a deductible do? The impact of cost-sharing on health care prices, quantities, and spending dynamics. *The Quarterly Journal of Economics*, 132(3):1261–1318.

Caves, D. W., Hanser, P., Herriges, J. A., and Windle, R. J. (1988). Load impact of interruptible and curtailable rate programs: Evidence from ten utilities. *IEEE Power Engineering Society Transactions*, 3(4):1757 – 1763.

Caves, D. W., Herriges, J. A., and Windle, R. J. (1992). The cost of electric power interruptions in the industrial sector: Estimates derived from interruptible service programs. *Land Economics*, 68(1):42–61.

Cena, K. and de Dear, R. (2001). Thermal comfort and behavioural strategies in office buildings located in a hot-arid climate. *Journal of Thermal Biology*, 26(4):409 – 414. International Thermal Physiology Symposium.

Chen, X. and Kleit, A. (2016). Money for nothing? Why FERC order 745 should have died. *Energy Journal*, 37(2):201 – 222.

Christiansen, M. R. (2016). FERC v. EPSA: Functionalism and the Electricity Industry of the Future. *Stanford Law Review Online 100*, 68.

Costa, D. L. and Kahn, M. E. (2013a). Do liberal home owners consume less electricity? A test of the voluntary restraint hypothesis. *Economics Letters*, 119(2):210 – 212.

Costa, D. L. and Kahn, M. E. (2013b). Energy conservation “nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment. *Journal of the European Economic Association*, 11(3):680–702.

Crawford, G. S., Pavanini, N., and Schivardi, F. (2018). Asymmetric information and imperfect competition in lending markets. *The American Economic Review*. (forthcoming).

Davis, L. (2012). *Evaluating the slow adoption of energy efficient investments: Are renters less likely to have energy efficient appliances?*, chapter 19. National Bureau of Economic Research Conference Report.

de Dear, R. and Brager, G. S. (1998). Developing an adaptive model of thermal comfort and preference. *ASHRAE Transactions*, 104:145–167.

Deryugina, T., MacKay, A., and Reif, J. (2019). The long-run elasticity of electricity demand: Evidence from municipal electric aggregation. *American Economic Journal: Applied Economics*, forthcoming.

Dewees, D. and Tombe, T. (2011). The impact of sub-metering on condominium electricity demand. *Canadian Public Policy / Analyse de Politiques*, 37(4):435–457.

EIA (2013). Heating and cooling no longer majority of U.S. home energy use. Technical report, United States Energy Information Administration (EIA).

EIA (2016). Carbon dioxide emissions coefficients by fuel. Technical report, United States Energy Information Administration (EIA).

EIA (2018). Electric power annual 2017. Technical report, United States Energy Information Administration (EIA).

Einav, L., Amy Finkelstein, S. R., Schrimpf, P., and Cullen, M. (2013). Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219.

Elinder, M., Escobar, S., and Petre, I. (2017). Consequences of a price incentive on free riding and electric energy consumption. *Proceedings of the National Academy of Sciences*, 114(4):3091–3096.

Eryilmaz, D., Smith, T. M., and Homans, F. R. (2017). Price responsiveness in electricity markets: Implications for demand response in the midwest. *Energy Journal*, 38(1):23–49.

Fanger, P. (1970). *Thermal comfort: Analysis and applications in environmental engineering*. Danish Technical Press.

Finkelstein, A., Gentzkow, M., and Williams, H. (2016). Sources of geographic variation in health care: Evidence from patient migration. *The Quarterly Journal of Economics*, 131(4):1681–1726.

Gillingham, K., Harding, M., and Rapson, D. (2012). Split incentives in residential energy consumption. *The Energy Journal*, 33(2).

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

Griefahn, B. and Künemund, C. (2001). The effects of gender, age, and fatigue on susceptibility to draft discomfort. *Journal of Thermal Biology*, 26(4):395 – 400. International Thermal Physiology Symposium.

Harding, J., Miceli, T. J., and Sirmans, C. (2000). Do owners take better care of their housing than renters? *Real Estate Economics*, 28(4):663–681.

He, J., Rejesus, R. M., Zheng, X., and Yorobe Jr, J. (2017a). Advantageous selection in crop insurance: Evidence from the philippines. *Journal of Agricultural Economics*.

He, J., Rejesus, R. M., Zheng, X., and Yorobe Jr, J. (2017b). Moral hazard and adverse selection effects of cost-of-production crop insurance: Evidence from the philippines. *Working paper*.

Hui, X., Saeedi, M., Shen, Z., and Sundaresan, N. (2016). Reputation and regulations: Evidence from eBay. *Management Science*, 62(12):3604–3616.

Indraganti, M. and Rao, K. D. (2010). Effect of age, gender, economic group and tenure on thermal comfort: A field study in residential buildings in hot and dry climate with seasonal variations. *Energy and Buildings*, 42(3):273 – 281.

Ito, K. (2014). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. *American Economic Review*, 104(2):537–63.

Jackson, C. K. and Schneider, H. S. (2011). Do social connections reduce moral hazard? Evidence from the New York City taxi industry. *American Economic Journal: Applied Economics*, 3(3):244–67.

Jackson, C. K. and Schneider, H. S. (2015). Checklists and worker behavior: A field experiment. *American Economic Journal: Applied Economics*, 7(4):136–68.

Jaffe, A. B. and Stavins, R. N. (1994). The energy paradox and the diffusion of conservation technology. *Resource and Energy Economics*, 16(2):91–122.

Jessoe, K., Papineau, M., and Rapson, D. (2018). Utilities included: Split incentives in commercial electricity contracts. Technical Report 029, E2e.

Jessoe, K. and Rapson, D. (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review*, 104(4):1417–38.

Karjalainen, S. (2007). Gender differences in thermal comfort and use of thermostats in everyday thermal environments. *Building and Environment*, 42(4):1594 – 1603.

Karjalainen, S. (2012). Thermal comfort and gender: A literature review. *Indoor Air*, 22(2):96–109.

Kema (2011). PJM empirical analysis of demand response baseline methods. Technical report, Kema, Inc.

Kingma, B. and van Marken Lichtenbelt, W. (2015). Energy consumption in buildings and female thermal demand. *Nature Climate Change*, 5:1054–1056.

Kox, M., van Eijk, L. T., Zwaag, J., van den Wildenberg, J., Sweep, F. C. G. J., van der Hoeven, J. G., and Pickkers, P. (2014). Voluntary activation of the sympathetic nervous system and attenuation of the innate immune response in humans. *Proceedings of the National Academy of Sciences*, 111(20):7379–7384.

Kuminoff, N. V., Smith, V. K., and Timmins, C. (2013). The new economics of equilibrium sorting and policy evaluation using housing markets. *Journal of Economic Literature*, 51(4):1007–1062.

Leffler, K. B. and Rucker, R. R. (1991). Transactions costs and the efficient organization of production: A study of timber-harvesting contracts. *Journal of Political Economy*, 99(5):1060–1087.

Levinson, A. and Niemann, S. (2004). Energy use by apartment tenants when landlords pay for utilities. *Resource and Energy Economics*, 26(1).

McAnany, J. (2018). 2017 Demand response operations markets activity report: April 2018. Technical report, PJM Demand Side Response Operations.

Monitoring Analytics (2018). State of the market report for pjm: January through september. Technical report, Monitoring Analytics, LLC.

Muzi, G., Abbritti, G., Accattoli, M. P., and dell'Omo, M. (1998). Prevalence of irritative symptoms in a nonproblem air-conditioned office building. *International Archives of Occupational and Environmental Health*, 71(6):372–378.

Myers, E. (2017). Are home buyers myopic? Evidence from capitalization of energy costs. Technical report, E2e.

Myers, E. (2018). Asymmetric information in residential rental markets: Implications for the energy efficiency gap. Technical report, E2e.

Nagashima, K., Yoda, T., Yagishita, T., Taniguchi, A., Hosono, T., and Kanosue, K. (2002). Thermal regulation and comfort during a mild-cold exposure in young Japanese women complaining of unusual coldness. *Journal of Applied Physiology*, 92(3):1029–1035.

Nakano, J., ichi Tanabe, S., and ichi Kimura, K. (2002). Differences in perception of indoor environment between Japanese and non-Japanese workers. *Energy and Buildings*, 34(6):615 – 621. Special Issue on Thermal Comfort Standards.

O'Connel, N., Pinson, P., Madsen, H., and O'Malley, M. (2014). Benefits and challenges of electrical demand response: A critical review. *Renewable and Sustainable Energy Reviews*, 39:686 – 699.

Papineau, M. (2017). Energy efficiency premiums in unlabeled office buildings. *The Energy Journal*, Volume 38(Number 4).

Parsons, K. C. (2002). The effects of gender, acclimation state, the opportunity to adjust clothing and physical disability on requirements for thermal comfort. *Energy and Buildings*, 34(6):593 – 599. Special Issue on Thermal Comfort Standards.

Pellerin, N. and Candas, V. (2003). Combined effects of temperature and noise on human discomfort. *Physiology and Behavior*, 78(1):99 – 106.

Pigou, A. C. (1920). *The Economics of Welfare*. Macmillan, London.

PJM (2013). 2012 Economic demand response performance report: Analysis of economic DR participation in the PJM wholesale energy market after the implementation of order 745. Technical report, PJM Interconnection.

PJM (2014a). Demand response and why it's important. Technical report, PJM Interconnection.

PJM (2014b). PJM economic demand resource in energy market. Technical report, PJM State and Member Training Department.

PJM (2017). Demand response. Technical report, PJM Interconnection.

Reiss, P. C. and White, M. W. (2005). Household electricity demand, revisited. *The Review of Economic Studies*, 72(3):853–883.

Reynolds, S. J., Black, D. W., Borin, S. S., Breuer, G., Burmeister, L. F., Fuortes, L. J., Smith, T. F., Stein, M. A., Subramanian, P., Thorne, P. S., and Whitten, P. (2001). Indoor environmental quality in six commercial office buildings in the midwest United States. *Applied Occupational and Environmental Hygiene*, 16(11):1065–1077. PMID: 11757903.

Rich, E. (2017). FERC order 825 - economic demand response. Technical report, PJM Market Settlements Subcommittee.

Ringsdorrf Jr., W. and Cheraskin, E. (1982). Vitamin C and tolerance of heat and cold: Human evidence. *Orthomolecular Psychiatry*, 11:128–131.

Schellen, L., Lichtenbelt, W. D. V. M., Loomans, M. G. L. C., Toftum, J., and Wit, M. H. D. (2010). Differences between young adults and elderly in thermal comfort, productivity, and thermal physiology in response to a moderate temperature drift and a steady-state condition. *Indoor Air*, 20(4):273–283.

Taylor, N. A., Allsopp, N. K., and Parkes, D. G. (1995). Preferred room temperature of young vs aged males: The influence of thermal sensation, thermal comfort, and affect. *The Journals of Gerontology: Series A*, 50A(4):M216–M221.

US Census (2013). American housing survey 1997-2013.

Veiga, A. and Weyl, E. G. (2016). Product design in selection markets. *The Quarterly Journal of Economics*, 131(2):1007–1056.

Walawalkar, R., Blumsack, S., Apt, J., and Fernands, S. (2008). An economic welfare analysis of demand response in the PJM electricity market. *Energy Policy*, 36:3692 – 3702.

Weisburd, S., Bird, D., and Ben-Porath, R. (2018). Adverse selection and moral hazard in the leasing market: Are buybacks the solution? Technical report, CEPR.

Wooldridge, J. (2010). *Econometric analysis of cross section and panel data*. MIT Press, Cambridge, Mass, 2 edition.

Young, A. J. (2010). *Homeostatic responses to prolonged cold exposure: Human cold acclimatization*, pages 419–438. John Wiley and Sons, Inc., supplement 14: handbook of physiology edition.