HOTSPOTS OF CLIMATE-DRIVEN INCREASES IN RESIDENTIAL ELECTRICITY DEMAND:

A Simulation Exercise Based on Household Level Billing Data for California

A White Paper from the California Energy Commission's California Climate Change Center

Prepared for: California Energy Commission

Prepared by: University of California, Berkeley

JULY 2012

CEC-500-2012-021

Maximilian Auffhammer
University of California, Berkeley

Anin Aroonruengsawat
Thammasat University

DISCLAIMER

This paper was prepared as the result of work sponsored by the California Energy Commission. It does not necessarily represent the views of the Energy Commission, its employees or the State of California. The Energy Commission, the State of California, its employees, contractors and subcontractors make no warrant, express or implied, and assume no legal liability for the information in this paper; nor does any party represent that the uses of this information will not infringe upon privately owned rights. This paper has not been approved or disapproved by the California Energy Commission nor has the California Energy Commission passed upon the accuracy or adequacy of the information in this paper.

ACKNOWLEDGEMENTS

We would like to thank the Public Interest Energy Research (PIER) Program at the California Energy Commission for generous funding of this work. Guido Franco has displayed a great deal of patience and provided useful suggestions throughout this project. We thank the University of California Energy Institute (UCEI) and the investor-owned utilities of California for letting us gain access to the billing data at the UCEI data center. We also thank Severin Borenstein and Koichiro Ito for helping us understand the electricity data. Alan Sanstad provided the electricity price scenarios. The Public Policy Institute of California provided the population projections at the county level. The downscaled global circulation model output was provided by Ed Maurer. Three anonymous referees have provided high quality comments, which have significantly improved this manuscript. We owe them our gratitude for looking at an early rough draft of this manuscript. All findings and remaining errors in this study are those of the authors.

ABSTRACT

One of the obvious modes of adaptation to higher temperatures due to climate change is the increased demand for cooling and decreased demand for heating in the built environment. California's residential sector uses relatively little electricity for heating, and it is therefore expected that the demand for electricity will increase as households operate existing air conditioners more frequently, and in many regions, will install air conditioners where there currently are few. This paper provides reduced form estimates of changes in electricity consumption due to increased use of installed cooling equipment under a hotter climate. This study adds to the literature by incorporating the change in temperature responsiveness due to likely increases in air conditioner penetration under climate change using a two-stage method. It shows that taking into account these capital investments may lead to higher projections of electricity consumption. These increases in projected electricity consumption were mapped to the ZIP codes in the study data. The paper shows suggestive evidence that more Caucasian and wealthy ZIP codes are projected to experience relatively smaller increases in consumption, while ZIP codes with a higher share of Latino population are projected to experience larger increases in consumption.

Keywords: climate change, vulnerability, electricity consumption, heating, cooling

Please use the following citation for this paper:

Auffhammer, Maximilian and Anin Aroonruengsawat (California Climate Change Center). 2012. ***Hotspots of Climate-Driven Increases in Residential Electricity Demand: A Simulation Exercise Based on Household Level Billing Data for California.*** California Energy Commission. Publication number: CEC-500-2012-021.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Unless otherwise noted, all tables and figures are provided by the author.

Section 1: Introduction

California is the world's eighth largest economy and home to 37 million people. While California's economy and population have grown by 2.83 percent and 1.54 percent per year between 1979 and 2010, its per capita electricity use has grown by only 0.25 percent per year (EIA SEDS 2012). This relatively constant per capita electricity consumption has been called the *Rosenfeld effect*, and is partially attributed to California's aggressive energy efficiency programs, which started in the early 1970s. Total electricity consumption for all sectors has grown by 6.4 percent over the most recent decade, with the residential sector being the fastest growing at 13 percent (EIA SEDS 2012).

While rising incomes and population have led to increases in the aggregate consumption of electricity, global climate change is expected to further increase this growth in the absence of additional policies. Warmer temperatures decrease the demand for heating, yet increase the demand for cooling in the built environment. Natural gas serves as the main fuel for heating in California and given current technology, electricity is the main source of energy for cooling. The increased demand for electricity due to a higher demand for cooling is expected to be significant —especially during peak times (e.g., Hayhoe et al. 2010; Miller et al. 2008). In previous work (Aroonruengsawat and Auffhammer 2012 [hereafter AA]), have shown that for the most likely Intergovernmental Panel on Climate Change (IPCC) emissions scenario, absent of population and income growth, additional policies, and changes in the adoption of air conditioners, residential electricity consumption may rise by between 1 to 6 percent by the end of the century. Most important, this study showed that electricity consumption is much more temperature sensitive during hot days in the Central Valley and southeastern parts of California than in the northern and coastal areas of the state. The previous study held the nonlinear relationship between electricity consumption to higher temperatures constant through time, and therefore only accounted for the intensive margin adjustment to higher temperatures (e.g., more frequent operation of existing air conditioners). It did not account for the extensive margin adjustments (e.g., the purchase of additional air conditioners in response to higher temperatures by California's households), which would change these response curves.

This current paper extends our previous work in three significant ways. First, we estimate the response function between electricity consumption and temperature at the five-digit ZIP code level instead of the California Energy Commission climate zone level. These response functions allow us to examine how the intensive margin adjustment ("increased usage of existing equipment") varies across 970 ZIP codes in our sample instead of only 16 climate zones. We use the same household level dataset used in AA, which contains the bills for all households served by California's investor-owned utilities for the years 2003–2006 matched to daily weather data. This first step is conceptually identical to our previous analysis, yet at a much more disaggregated level.

Second, we explain cross-sectional variation in these "first stage" estimated slopes of each ZIP code's temperature response function as a function of socio-demographics and "climate." Technically speaking, in this "second stage" we regress the slope of each ZIP code's temperature response function in different temperature bins on observable characteristics of the population across ZIP codes and an average of summer temperatures, which we call (summer) *climate.* We therefore separate the impact of socio-demographics (e.g., income, ethnicity) on temperature response from the direct effect of climate (average JJA[1] temperatures). The estimated marginal effect of climate on the slope of the response function allows us to capture extensive margin adjustments to long run changes in climate. We use downscaled predictions from three well-known climate models to simulate future household electricity consumption at the ZIP code level under climate change, taking into account both intensive ("first stage") and extensive margin ("second stage") adjustments.

Finally, we calculate the above projections with and without projections of population growth. We examine the empirical distribution of these estimated impacts along socio-demographic lines. Taking the 2000 Census distribution of ethnic groups and incomes as given, we study whether certain groups are especially vulnerable to increases in the electricity demand from climate change.

Before we proceed with the detailed discussion of our approach and related literature, it is important to state a few caveats. These projections do not take into account policy-induced or exogenous improvements in the energy efficiency of air conditioners and buildings or changes in the size of buildings in response to changes in climate (either voluntary or due to urban planning). They therefore clearly represent an upper bound. Further, while we interpret simulated increases in electricity to be due to cooling demand, it is quite possible that individuals spend more time inside engaging in electricity-intensive activities partially or wholly unrelated to cooling (e.g., television watching, lighting). If our econometric estimates are consistent, our scenarios therefore present an upper bound of the impacts of climate change conditional on the projections of climate and population. If the stated policy goal is to reach significant reductions in greenhouse gas emissions, these scenarios can serve as a baseline scenario. This baseline scenario does not contain changes from additional policy, autonomous improvements in energy efficiency of air conditioners, and building size.

The remainder of the paper is organized as follows: Section 2 briefly reviews the literature assessing the impacts of climate change on electricity demand, as well as the economics literature on air conditioner adoption. Section 3 describes the sources of the data used in this study. Section 4 contains the econometric model, and Section 5 has the estimation results. Section 6 discusses electricity consumption increases due to climate and population change, and Section 7 concludes.

[1] June, July, August

Section 2: Literature Review

The literature on climate change impacts estimation can be divided into two approaches. In the engineering literature, large-scale bottom-up simulation models are utilized to simulate future electricity demand under varying climate scenarios. The advantage of the simulation model approach is that it allows one to simulate the effects of climate change given a wide variety of technological and policy responses. The drawback to these models is that they contain a large number of response coefficients and make a large number of assumptions about the evolution of the capital stock, for either of which there is little empirical guidance.

The earliest impacts papers adopt this simulation approach and suggest that global warming will significantly increase energy consumption. Cline (1992) provides the earliest study on the impacts of climate change in his seminal book, *The Economics of Climate Change.* His section dealing with the impact on space cooling and heating relies on an earlier report by the U.S. Environmental Protection Agency (1989). That study of the potential impact of climate change on the United States uses a utility planning model developed by Linder et al. (1987) to simulate the impact on electric utilities in the United States and finds that increases in annual temperatures ranging from 1.0°C–1.4°C (1.8°F–2.5°F) in 2010 would result in demand of 9 percent to 19 percent above estimated new capacity requirements (peak load and base load) in the absence of climate change. The estimated impacts rise to 14 percent and 23 percent for the year 2055 and an estimated 3.7°C (6.7°F) temperature increase.

Baxter and Calandri (1992) provide another early study in this literature and focus on California's electricity use. In their study they utilize a partial equilibrium model of the residential, commercial, agriculture, and water pumping sectors to examine total consumption, as well as peak demand. They project electricity demand for these sectors to the year 2010 under two global warming scenarios: a rise in average annual temperature of 0.6°C (1.1°F) (Low scenario) and of 1.9°C (3.4°F) (High scenario). They find that electricity use increases from the constant climate scenario by 0.6 percent to 2.6 percent, while peak demand increases from the baseline scenario by 1.8 percent to 3.7 percent.

Rosenthal et al. (1995) focus on the impact of global warming on energy expenditures for space heating and cooling in residential and commercial buildings. They estimate that a 1°C (1.8°F) increase in temperature will *reduce* U.S. energy expenditures in 2010 by $5.5 billion (1991 dollars).

The more recent economics literature has frequently applied the statistics-based econometric approach to impacts estimation, which is the approach we adopt in the current study. While there is a large literature on econometric estimation of electricity demand, the literature on climate change impacts estimation is small and relies on panel estimation of heavily aggregated data or cross-sectional analysis of more micro-level data. The first set of papers attempts to explain variation in a cross section of energy expenditures based on survey data to estimate the impact of climate change on fuel consumption choices. Mansur et al. (2008) and Mendelsohn

(2003) endogenize fuel choice, which is usually assumed to be exogenous. They find that warming will result in fuel switching towards electricity. The drawback of the cross-sectional approach is that one cannot econometrically control for unobservable differences across firms and households, which may be correlated with weather/climate. If that is the case, the coefficients on the weather variables and corresponding impacts estimates may be biased.

Instead of looking at a cross section of firms or households, Franco and Sanstad (2008) explain pure time series variation in hourly electricity load at the grid level over the course of a year. They use data reported by the California Independent System Operator for 2004 and regress them on average daily temperature. The estimates show a nonlinear impact of average temperature on electricity load and a linear impact of maximum temperature on peak demand. They link the econometric model to climate model output from three different general circulation models (GCMs) forced using three IPCC scenarios (A1Fi, A2, and B1) to simulate the increase in annual electricity and peak load from 2005–2099. Relative to the 1961–1990 base period, the range of increases in electricity and peak load demands are 0.9 to 20.3 percent and 1.0 to 19.3 percent, respectively. Crowley and Joutz (2003) use a similar approach where they estimate the impact of temperature on electricity load using hourly data in the Pennsylvania, New Jersey, and Maryland Interconnection. Some key differences, however, are that they control for time-fixed effects and define the temperature variable in terms of heating and cooling degree days. They find that a 2°C (3.6°F) increase in temperature results in an increase in energy consumption of 3.8 percent of actual consumption, which is similar to the impact estimated by Baxter and Calandri (1992).

Deschênes and Greenstone (2011) provide the first panel data-based approach to estimating the impacts of climate change on residential electricity demand. They explain variation in U.S. state-level annual panel data of residential energy consumption using flexible functional forms of daily mean temperatures. The identification strategy behind their paper, which is one we will adopt here as well, relies on random fluctuations in weather to identify climate effects on electricity demand. The model includes state fixed effects, census division by year fixed effects, and controls for precipitation, population, and income. The temperature data enter the model as the number of days in 20 predetermined temperature intervals. The authors find a U-shaped response function where energy consumption is higher on colder and hotter days. The impact of climate change on annual residential energy consumption for the Pacific Census Region (California, Oregon, and Washington) by 2099 is approximately nine percent—yet not statistically different from zero. They are careful to point out that their estimates likely present an upper bound. The panel data approach allows one to control for differences in unobservables across the units of observation, resulting in consistent estimates of the coefficients on temperature.

Aroonruengsawat and Auffhammer (2012), which corrects an error in previous work, is the first paper using a panel of household level electricity billing data to examine the impact of climate change on residential electricity consumption. They identify the effect of temperature on electricity demand using within-household variation in temperature, which is made possible through variation in start dates and lengths of household billing periods. Since their dataset is a panel, they can control for household fixed effects, month fixed effects, and year fixed effects.

The literature looking at extensive margin adjustments, not necessarily just to higher temperatures, provides significant insight into how households decide to invest in durable goods such as air conditioners and decide how intensively to use them. Biddle (2008) studies the spread of air conditioning in the United States in the post World War II economy. In the early 1950s air conditioners were mainly found in public spaces (e.g., movie theatres and supermarkets). He notes that in 1955 the residential air conditioner penetration in the United States was below 2 percent nationally. A quarter of a century later that fraction had risen to 50 percent, with half of those households having installed central air conditioning units. There was significant heterogeneity in the penetration, where half of the residences in the Northeast were air-conditioned, and some urban areas in Texas and Florida had penetration rates in excess of 90 percent.

He argues that on the one hand, the strong downward trend in prices of air conditioners is a potential driver of the increased adoption of this technology. Further, electricity prices dropped significantly during the 1950s and 1960s and then rose again during the 1970s. During this entire period incomes rose substantially, which suggests that the falling costs of installation and operation, combined with rising incomes, drove the adoption of air conditioners during this period. To determine the relative importance of these factors, Biddle (2008) matches the air conditioning indicators with the corresponding socioeconomic characteristics from three Census cross sections for 1960, 1970, and 1980 to electricity rates, incomes, and detailed climate variables (e.g., cooling and heating degree days, wind speed, relative humidity). He uses a reduced form econometric model, which accounts for changes in incomes, prices, and weather to explain the heterogeneity in penetration.

His estimated income elasticities are positive in most specifications and years, even though they vary widely in size and significance across samples and estimation techniques. The price elasticities are generally negative and often significant, although they vary significantly in magnitude as well. Biddle (2008) concludes that while rising incomes and dropping real prices of electricity drove the adoption of air conditioners in the United States during the two decades he studies, the changing housing stock had an impact on the costs of adopting air conditioning, which accelerated these trends. He shows weaker evidence that significant drops in the initial cost of installing air conditioners also promoted the adoption of these units. He is careful to point out that these results may be confounded by improved efficiency of the air conditioning units and changes in the housing stock, for which he cannot control.

Sailor and Pavlova (2003) use data on air conditioning penetration for 39 U.S. cities to parameterize a relationship between cooling degree days (CDDs) and market saturation. They take issue with existing estimates that electricity consumption rises by 2 to 4 percent for each degree Celsius in warming. They show penetration data from the American housing survey for 39 cities for the year 1994–1996 for both central and window units. There is clear nonlinear relationship between CDDs and penetration. Further, a significant number of cities have air conditioning penetration below 80 percent, suggesting that there are two margins of adjustment under climate change: Increased adoption of air conditioners and increased usage. Ignoring the adoption decision would lead to an underestimation of future electricity consumption. They estimate a relationship between saturation and CDDs, which uses the notion that current hotter cities have patterns of adoption of cities that are currently cooler, but are expected to be hotter under climate change. The authors show that there is a relationship between the two variables. In their analysis of electricity consumption the authors use state per capita electricity consumption as a proxy for city-level per capita consumption, and their statistical model allows both for increased usage of equipment due to more CDDs and increased adoption of air conditioners due to higher CDDs. They show that simulated increases are significant, and also indicate that taking into account the adoption decision matters in making forecasts. Sailor and Pavlova (2003) note that "Based on these results, Los Angeles' per capita residential electricity consumption is projected to increase by 8% in July for a 20% increase in CDD. If the market saturation were assumed to remain constant, however, the projection would be for only a 5% increase."

Rapson (2011) estimates a dynamic structural model of air conditioner adoption and simulates the impacts of a carbon tax and energy-efficiency standards. The three relevant findings from his study are that higher energy efficiency of air conditioners raises demand for these units and decreases overall consumption. Second, higher electricity prices have a negative effect on current consumption and a marginal negative effect on adoption. Finally, lower prices of air conditioners lead to higher demand and marginally lower consumption. He uses five cross sections of the Energy Information Administration's Residential Energy Consumption Survey (RECS), which he matches to air conditioner prices and efficiencies.

A very recent report by the Energy Information Administration (2011) for the most recent round of RECS shows further growth in air conditioner (AC) penetration on the United States. There is little slowdown in the growth of air conditioner penetration. Eighty-seven percent of U.S. households had air conditioning in 2009, which is the latest year of data. The EIA (2011) notes that "Wider use has coincided with much improved energy efficiency standards for AC equipment, a population shift to hotter and more humid regions, and a housing boom during which average housing sizes increased."

Central AC units are most common in the "South" Census Region, and window units are most common in the "NorthEast" Census region. There is little variation in usage over the summer, as the percentage of households using AC during the summer is between 30 and 40 percent—with the exception of the South, where 67 percent of households run their air conditioners all summer. Further, newer homes are most likely to have central AC; whereas, older homes are more likely to have no air conditioning or window units, as retrofitting with central AC has non-trivial transactions costs. The Energy Information Administration (2011) further notes that there is significant heterogeneity in the penetration and type of AC units installed across the income spectrum.

McNeil and Letschert (2008, 2010) provide a comprehensive model of adoption of air conditioners and appliances using cross-country data. They collect a cross section of data on adoption rates by country from a number of micro-level survey studies and IEA data. They acknowledge that the utility of owning an air conditioner is climate dependent. Further, they point out that the cost of owning an air conditioner is relatively high. In order to break the correlation between income and temperature in their dataset, the authors take a two-step approach. They follow Sailor and Pavlova (2003) and estimate a relationship between saturation and cooling degree days for 39 U.S. cities. They then assume that for a given CDD, developing country air conditioner saturation will approach this frontier, but never exceed it. They assume that diffusion is a function of the climate maximum calculated above and income. They then estimate a relationship between income and air conditioner saturation, which non-surprisingly shows that income is a major driver of the speed of adoption.

The current paper adds to these two literatures more broadly by examining how residential electricity consumption changes at high temperature at a highly disaggregated level (ZIP code) using observed household level consumption data over 48 billing periods. It then goes one step further and examines the relative importance of climate and sociodemographics (e.g., income) in influencing the temperature sensitivity of electricity consumption during hot days.

Section 3: Data

3.1 Residential Billing Data

The University of California Energy Institute, jointly with California's investor-owned utilities, established a confidential data center, which contains the complete billing history for all households serviced by Pacific Gas and Electric, Southern California Edison, and San Diego Gas and Electric. We use data for these three utilities for the years 2003–2006. While we have data on additional years for some utilities, due to the discrepancy of the time series coverage of billing data from each utility, we limit the study to the years 2003–2006, where billing data from all three utilities are avaliable.

The dataset contains the complete consumption and expenditure information for each residential customer's bills over this four-year period. Specifically, we observe an identity for the physcial location, a service account number, bill start-date, bill end-date, total electricity consumption (in kilowatt-hours, kWh), and the total amount of the bill (in $) for each billing cycle, as well as the five-digit ZIP code of the premises.[2] Only customers who were individually metered are included in the dataset. We cannot reliably identify households who have moved and therefore refrain from using this as a source of econometric identification. For the purpose of this paper, a *customer* is defined as a unique combination of premise and service account number. It is important to note that each billing cycle does not follow the calendar month, and the length of the billing cycle varies across households, with the vast majority of households being billed on a 25–35 day cycle. Hereafter, this dataset is referred to as "billing data."

The billing dataset contains 300 million observations, which exceeds our ability to conduct estimations using standard statistical software. We therefore resort to sampling from the population of residential households to conduct econometric estimation. We designed the following sampling strategy, which differs from AA. First we only sample from households with regular billing cycles, namely 25–35 days in each billing cycle and which have a complete set of 47 or 48 bills over the period of 2003–2006.[3] We also remove bills with average daily consumption less than 2 kWh or more than 80 kWh, since we are concerned that these outliers are not residential homes, but rather vacation homes and small-scale manufacturing facilities. Our data do not contain multi-family single-meter homes, so these homes are not represented in the study. Our results should be interpreted keeping this in mind when judging the external validity of our results.

There is significant variation in bill level consumption across and within households. Because across-household variation may be driven by unobservable characteristics at the household

---

[2] The premise identification number does not change with the occupant of the residence. The service account number, however, changes with the occupant of the residence. This information is subject to sizable measurement error, so we do not make use of "movers" for econometric identification purposes.

[3] With the regular billing cycle, there should be about 48 bills for the existing households during 2003 to 2006. AA did not require a minimum of 47 bills per household.

level (e.g., income, physcial building characteristics, and installed capital), we will control for unobservable confounders at the household level using fixed effects and use bill-to-bill variation at the household level as our source of identifying variation. To prceed with estimation at the ZIP code level, from the population subject to the restrictions above, we take a random sample from each ZIP code of up to 1,000 households, resulting in up to 48,000 bills per ZIP code. We drop ZIP codes with fewer than 100 households with a full billing record. Figure 1 displays the ZIP codes for which we have data, which cover a large portion of the state and represent approximately 80 percent of California's population.

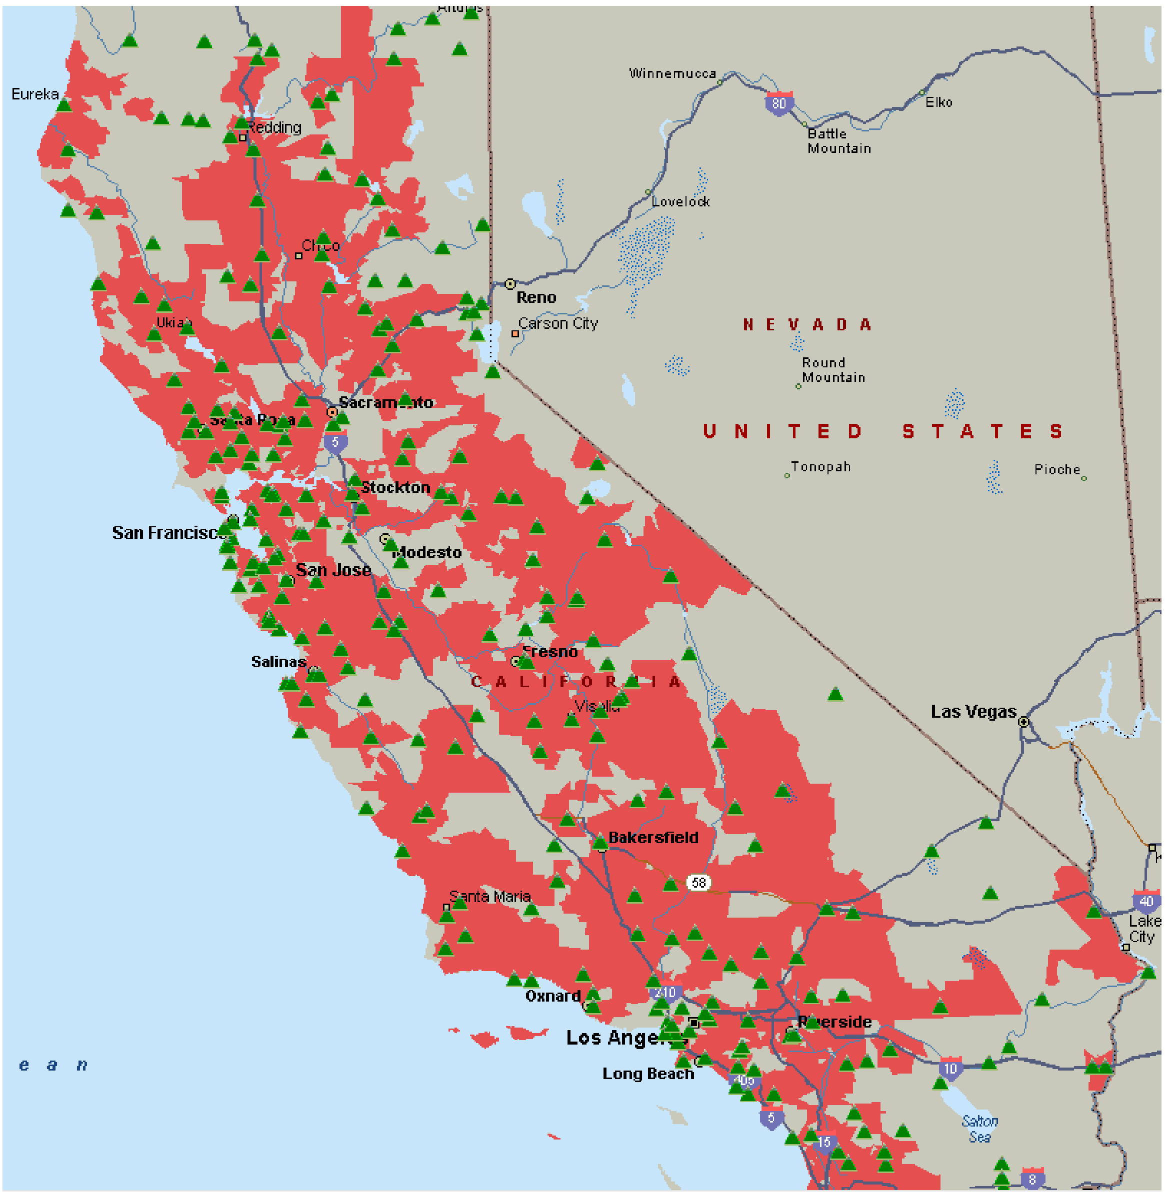

Notes: The map above displays the five-digit ZIP codes for which we have households with a record of at least 47 bills during 2003–2006. This sample differs from the sample used by AA. The green triangles display the location of NOAA weather stations satisfying the criteria set out in Section 3.2.

Figure 1: ZIP Codes with Sampled Residential Electricity Consumption (red) and National Oceanic and Atmospheric Administration (NOAA) Cooperative Weather Stations (green)

While we have at least 47 bills for 1,000 households for about 25 percent of the ZIP codes in our sample, we have smaller samples for the remaining ZIP codes, mostly due to the sparse population levels in some of the northern and interior counties. We plot the distribution of the number of households across ZIP codes in Figure 2. No single ZIP code is responsible for more than 0.5 percent of total consumption.

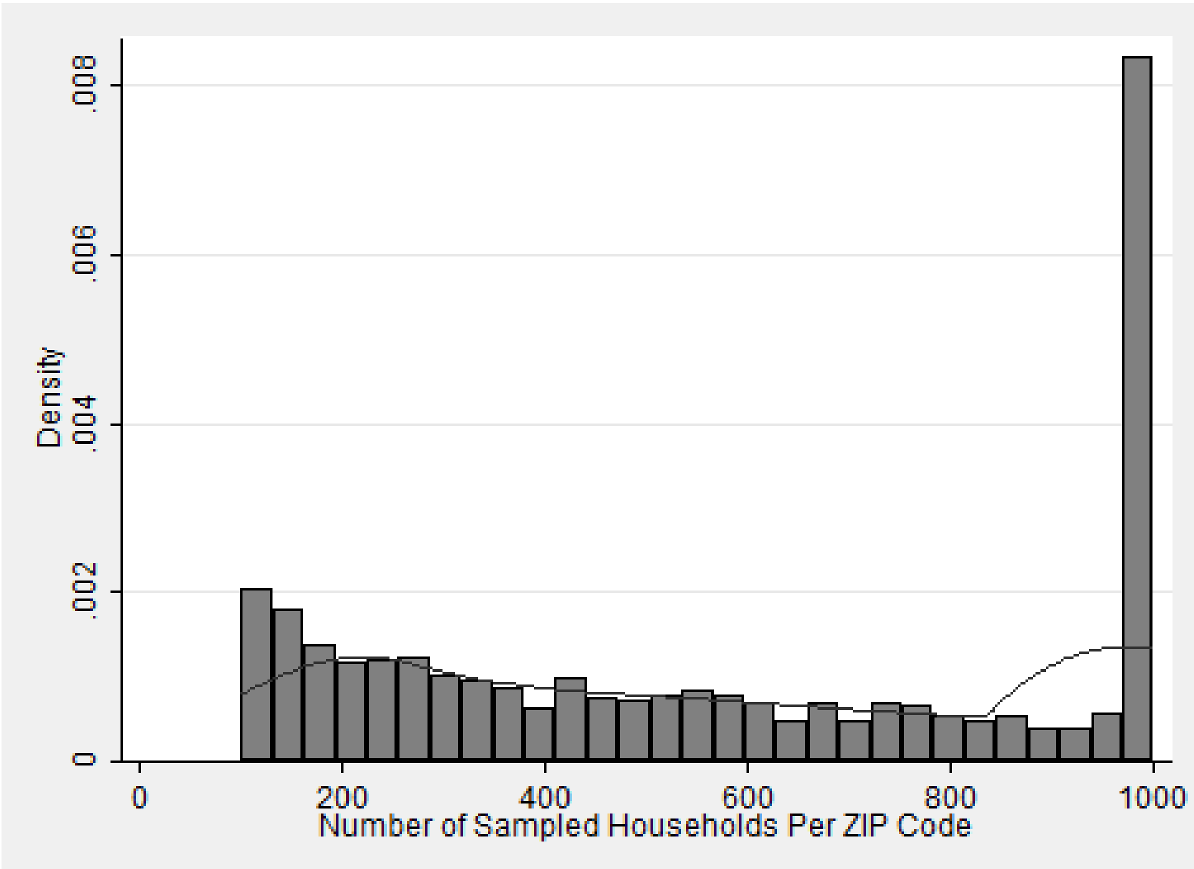

Figure 2: Distribution of Number of Sampled Households across Five-Digit ZIP Codes

Since summarizing consumption data for the 970 ZIP codes that ultimately ended up in our estimation cannot be meaningfully presented here, we summarize the consumption, pricing, and weather data for our sample by California Energy Commission climate zone in Table 1. There is great variability in average usage across climate zones, with the highest consumption levels in the hot Central Valley and China Lake (Southeast) zones. The average electricity price paid is similar across zones for our sample, ranging between 13 and 15 cents per kilowatt-hour.

Table 1: Summary Statistics

CEC Climate Zone	Bills	Usage Per Bill (kWh) Mean	s.d.	Average Price ($/kWh) Mean	s.d.	Percentage of Days in Percentiles of overall Temperature Distribution: Bottom 5%	Bottom 10%	Top 10%	Top 5%
1: Arcata	633,585	572	339	0.13	0.03	1.14%	4.88%	1.08%	0.16%
2: Santa Rosa	1,343,520	616	359	0.13	0.02	0.41%	2.76%	2.83%	0.58%
3: Oakland	2,434,224	479	291	0.13	0.02	0.00%	0.35%	1.11%	0.17%
4: Sunnyvale	2,051,342	603	337	0.13	0.02	0.06%	0.81%	2.52%	0.43%
5: Santa Maria	862,714	506	304	0.13	0.03	0.07%	1.43%	2.70%	0.26%
6: Los Angeles	2,099,911	530	315	0.13	0.03	0.02%	0.02%	2.35%	0.41%
7: San Diego	2,104,583	503	303	0.15	0.03	0.00%	0.00%	0.90%	0.08%
8: El Toro	1,789,992	577	331	0.14	0.03	0.00%	0.02%	5.72%	1.32%
9: Pasadena	2,315,756	614	353	0.13	0.02	0.12%	0.31%	11.40%	3.32%
10: Riverside	2,203,563	674	374	0.14	0.03	1.55%	3.15%	9.71%	3.21%
11: Red Bluff	1,684,772	761	400	0.13	0.02	3.43%	11.16%	12.70%	4.95%
12: Sacramento	2,088,723	696	374	0.13	0.02	0.34%	2.55%	10.48%	3.40%
13: Fresno	1,876,201	727	409	0.13	0.02	0.25%	2.72%	25.09%	13.24%
14: China Lake	1,316,191	687	375	0.13	0.03	3.94%	9.58%	22.48%	13.50%
15: El Centro	487,142	674	465	0.13	0.03	0.59%	1.84%	22.79%	14.83%
16: Mount Shasta	1,470,624	602	369	0.13	0.02	8.43%	15.11%	12.11%	5.27%

Note: The table displays summary statistics for residential electricity consumption for the sample used in the estimation. s.d. = standard deviation.

3.2 Weather Data

To generate daily weather observation to be matched with the household electricity consumption data, we use the Cooperative Station Dataset published by National Oceanic and Atmospheric Administration's (NOAA's) National Climate Data Center (NCDC). The dataset contains daily observations from more than 20,000 cooperative weather stations in the United States, U.S. Caribbean Islands, U.S. Pacific Islands, and Puerto Rico. Data coverage varies by station. Since our electricity data cover California for the years 2003–2006, the dataset contains 370 weather stations reporting daily data. In the dataset we observe daily minimum and maximum temperature, as well as total daily precipitation and snowfall. Since the closest meaningful geographic identifier of our households is the five-digit postal ZIP code, we select stations as follows. First, we exclude any stations not reporting data in all years. Further we exclude stations reporting fewer than 300 observations in any single year and stations at elevations more than 7,000 feet above sea level, which leaves us with 274 "valid" weather stations.[4] Figure 1 displays the distribution of these weather stations across the state (as green triangles). While there is good geographic coverage of weather stations for our sample, we do

[4] The cutoff of 300 valid days is arbitrary. If we limit the set of weather stations to the ones providing a complete record, we would lose roughly half of all stations. We conducted robustness checks using different cutoff numbers, and the estimation results are robust.

not have a weather station reporting data for each ZIP code. To assign a daily value for temperature and rainfall, we need to assign a weather station to each ZIP code. We calculate the Vincenty distance (which takes into account the curvature of the earth) of a ZIP code's centroid to all valid weather stations and assign the closest weather station to that ZIP code. As a consequence of this procedure, each weather station on average provides data for approximately ten ZIP codes.

Since we do not observe daily electricity consumption by household, but rather monthly bills for billing periods of differing length, we required a complete set of daily weather observations. The NCDC data have a number of missing values, which we fill in using the following algorithm. If a station is missing values for minimum/maximum temperature or precipitation, we regress the weather outcome of interest on the same variable for the ten closest stations each reporting data for at least 200 days a year. We then use the predicted value from this regression to fill in the missing observation. If there still are missing values, due to incomplete time series from the ten neighboring stations, we regress the series on the nine closest stations and use the predicted values. We repeat this process until all missing values for the station are filled. We end up with a complete set of time series for minimum temperature, maximum temperature, and precipitation for the 274 weather stations in our sample. To ensure that we are not fabricating bad quality data, we set aside 10 percent of the observed data and made these calculations without these data. We run our algorithm on these artificially incomplete series. When regressing the resulting series for temperature on the actual series, the intercept of the regression is indistinguishable from zero and the slope coefficient is indistinguishable from 1. The correlation coefficient for the two series is 0.99. For the remainder of our empirical analysis, we use these patched series as our observations of weather.[5]

3.3 Other Data

In addition to quantity consumed and average bill amount, unfortunately all we know about the household is the five-digit ZIP code in which it is located. We purchased socio-demographics at the ZIP code level from a firm aggregating this information from census estimates (zip-codes.com). We only observe these data for a single year (2006). There are 1,671 five-digit ZIP codes, which can be assigned to polygons via a geographical information system (GIS). Our sample contains households for 970 of these ZIP codes. We do not have sufficient data for households in the remaining 701 ZIP codes. These remaining ZIP codes are either not served by the three utilities, or we do not have households with a sufficient number of bills for them. Table 2 shows summary statistics for both the ZIP codes in our sample and the ZIP codes for which we do not have electricity data. The ZIP codes in our sample represent 77.09 percent of California's population. The ZIP codes in our sample are younger, richer, have more expensive homes, have slightly more persons per household, and have a lower proportion of Caucasians and a higher proportion of African Americans and Asians. There seems to be no statistically significant difference in mean summer weather, which is important for our

[5] Inverse distance weighting provides an alternate approach to filling in missing temperature values, which given the good fit of our algorithm we have not explored further in this study.

purposes. Taking these differences into consideration is again important when judging the external validity of our estimation and simulation results.

We will not make explicit use of this information in our first-stage regression, but control for the observable sources of variation in our cross-sectional second stage, which does not allow for a fixed effects strategy by design. The variables we will make use of in the second stage are population, income, shares of ethnic groups, and climate. We will also make use of these variables when we break down the ZIP code impacts to identify vulnerable groups and areas in our final step.

Table 2: Summary Statistics for ZIP Codes In and Out of Sample

Variable	In Sample			Not In Sample			
	n	Mean	Std. Dev.	Obs	Mean	Std. Dev.	t |p-value|
Population ('000)	970	26.91	20.49	701	11.07	17.15	<0.001
Share White	970	0.68	0.21	701	0.72	0.22	0.002
Share Afr. Am.	970	0.05	0.09	701	0.04	0.08	0.001
Share Hispanic	970	0.24	0.21	701	0.24	0.25	0.995
Share Asian	970	0.09	0.11	701	0.05	0.08	<0.001
Share Hawaiian	970	0.00	0.00	701	0.00	0.00	<0.001
Share Nat. American	970	0.01	0.02	701	0.02	0.05	<0.001
Share Male	970	0.50	0.03	701	0.52	0.06	<0.001
Person per Household	970	2.82	0.57	701	2.69	0.67	<0.001
Mean Home Value (10kUS$)	970	25.29	17.57	701	19.68	17.61	<0.001
Mean HH Income (1kUS$)	970	51.82	21.48	701	40.20	19.77	<0.001
Median Age	970	36.25	6.67	701	37.69	8.41	<0.001
Elevation	970	388.87	643.83	701	727.73	1176.10	<0.001
JJA Average Temp 1980-1999	970	20.82	3.29	367	20.93	3.89	0.638

Note: This table displays the differences in observable characteristics between ZIP codes that we observed at least 100 household's complete bills for the period 2003–2006 and ZIP codes that we did not. HH = household.

Section 4: Econometric Estimation Strategy

4.1. First Stage: The Usage Response to Temperature

As discussed in the previous section, we observe each household's monthly electricity bill for the period 2003–2006. Equation 1 below shows our main estimating equation, which is a simple log-linear equation, which has commonly been employed in aggregate electricity demand estimation and climate change impacts estimation (e.g., Deschênes and Greenstone 2011).

$$\ln(q_{it}) = \sum_{p=1}^{k} \beta_p D_{pit} + \gamma Z_{it} + \alpha_i + \phi_m + \varphi_y + \varepsilon_{it} \qquad (1)$$

where $\ln(q_{it})$ is the natural logarithm of household *i*'s electricity consumed in kilowatt-hours during billing period *t*. D_{pit} are our measures of temperature, which we discuss in detail below. Z_{it} are observed confounders at the household level, α_i are time invariant household fixed effects, ϕ_m are month of year fixed effects, and φ_y are year fixed effects. ε is a stochastic error term.

For estimation purposes, our unit of observation is a unique combination of premise and service account number, which is associated with an individual *and* structure. We thereby avoided the issue of having individuals moving to different structures with more or less efficient capital or residents with different preferences over electricity consumption moving in and out of a given structure.

California's housing stock varies greatly in its energy efficiency and installed energy-consuming capital. Further, California's population is not randomly distributed across ZIP codes. We suspect that there may be differences in attitudes towards cooling, installed capital, quality of construction across ZIP codes, and the associated demographics and capital. We estimated Equation 1 separately for each of the 970 ZIP codes discussed in the data section, which are also displayed in Figure 1. The motivation for doing so is that we would expect the relationship between consumption and temperature to vary across these ZIP codes.

The main variables of interest in this paper are those measuring temperature. Following recent trends in the literature and consistent with AA we include our temperature variables in a way that imposes a minimal number of functional form restrictions in order to capture potentially important nonlinearities of the outcome of interest—electricity consumption—in weather (e.g., Schlenker and Roberts 2006; Deschênes and Greenstone 2011). We achieve this by sorting each day's mean temperature experienced by household *i* into one of *k* temperature bins.[6] For the

[6] We use mean daily temperature as our temperature measure. This allows a flexible functional form in a single variable. An alternate strategy we will explore in future work is separating the temperature variables into minimum and maximum temperature, which are highly correlated with our mean temperature measure.

purposes of this study, we use the same set of bins for each ZIP code in the state. In order to define a set of temperature bins we split the state's temperature distribution into a set of percentiles and use those as the bins sorting. Aroonruengsawat and Auffhammer (2012) show that the alternative approach of using equidistant five-degree bins yields almost identical results. We therefore refrain from following this approach, due to the computational intensity involved. As a result, not each ZIP code will have observations in each bin. The northern ZIP codes, for example, do not experience days in the upper bins, while the southwestern parts of California have few days in the coldest bins.

We split the temperature distribution into deciles, yet break down the upper and bottom decile further to include buckets for the first, fifth, ninety-fifth, and ninety-ninth percentile to account for extreme cold/ heat days. We therefore have a set of 14 buckets which we use for each household, independent of in which climate zone the household is located.[7] For each household, we count the number of days the mean daily temperature falls into each bin and record this as D_{pit}. The main coefficients of interest to the later simulation exercise are the β_p coefficients, which measure the impact of one more day with a mean temperature falling into bin p on the log of household electricity consumption. For small values, β_p's interpretation is approximately the percent increase in household electricity consumption during a billing period from experiencing one additional day in that temperature bin.

Table 1 displays the heterogeneity in the weather distribution across climate zones. The Central Valley and the southeastern areas of the state for example, are significantly hotter on average, as shown by the larger share of days spent in the top fifth percentile of the overall temperature distribution.

Z_{it} is a vector of observable confounding variables, which vary across billing periods and households. The first of two major confounders that we observe at the household level are the average electricity price for each household for a given billing period. California utilities price residential electricity on a block rate structure. The average price experienced by each household in a given period is therefore not exogenous, since marginal price depends on consumption (q_{it}). Identifying the price elasticity of demand in this setting is problematic, and a variety of approaches have been proposed (e.g., Hanemann 1984; Reiss and White 2005). The maximum likelihood approaches are computationally intensive and given our sample size cannot be feasibly implemented here. We have run our models by including price directly, instrumenting for it using lagged prices and omitting it from estimation. The estimation results are almost identical for all three approaches, which is reassuring. While one could tell a story that higher temperatures lead to higher consumption, and therefore higher marginal prices for some households, this bias seems to be negligible given our estimation results. In the estimation and simulation results presented in this paper, we omit the average price from our main regression. The second major time-varying confounder is precipitation in the form of rainfall.

---

[7] The cutoffs for the bins are 28.5, 38.5, 43, 49, 52.5, 56, 59.5, 63, 66.5, 71, 77, 82, and 91.5 degrees Fahrenheit mean daily temperature.

We calculate the amount of total rainfall for each of the 274 weather stations, filling in missing values using the same algorithm discussed in the previous section. We control for rainfall using a second-order polynomial in all regressions.

To credibly identify the effects of temperature on the log of electricity consumption, we require that the residuals conditional on all right-hand side variables be orthogonal to the temperature variables, which can be expressed as $E[\varepsilon_{it} D_{pit} \mid D_{-pit}, Z_{it}, \alpha_i, \phi_m, \varphi_y] = 0$. Since we control for household fixed effects, identification comes from within household variation in daily temperature after controlling for confounders common to all households (e.g., business cycle effects), rainfall, and average prices. We estimate Equation 1 for each of the 970 ZIP codes in our sample using a least-squares fitting criterion and a clustered variance covariance matrix. This approach serves as the first stage in our overall methodology and matches the approach adopted by AA. We must make the assumption that response to slowly changing climate over this four-year period is small in order to be able to interpret our coefficients as the intensive margin adjustment—the changes in usage of existing equipment in response to changing temperature.

4.2. Second Stage: The Long-Run Response to Temperature

The second long-run mechanism of adaptation will likely be the adoption of additional air conditioners in existing homes and new construction. One can easily imagine that if La Jolla's future climate resembles that of current day Fresno during the summers, the wealthy residents of La Jolla will install (additional) cooling equipment in their homes. We provide an attempt to quantify the magnitude of this response. We estimate equations of the following form:

$$\beta_i^p = \delta_1 + \delta_2 C_i + \delta_3 Z_i + \eta_i \tag{2}$$

where β_i^p is a measure of ZIP code *i*'s temperature responsiveness in bin *p*. We run 14 different regressions, one for each temperature bin, where β_i^p is ZIP code *i*'s slope of the temperature response function for bin $p \in [1;14]$.

The variable C_i in Equation 2 is summer time average temperature during the months June, July, and August during the years 2003–2006 from the downscaled National Center for Atmospheric Research (NCAR) parallel climate model (PCM1) data for scenario A2. Using the downscaled model output makes a more straightforward exercise out of sample predictions since concerns about local bias to the weather station data is not a concern by design. As the climate model output is downscaled using weather station data, this should serve as a good approximation of actual climate for each ZIP code *i*. The variable(s) Z_i are any confounders that may affect the temperature response of the population in ZIP code *i*. The confounders we consider here are income, as higher-income households are thought to more easily afford the capital expenditure of an air conditioner and its associated operating expense (Rapson 2011). Further, one could argue that different population groups may have differential preferences for temperature. We therefore control for shares of Caucasian, Latino, and African American population. We leave out the groups "Asian" and "Other" as defined by the Census as a baseline. We further control for population and elevation in a ZIP code. While we will not use these estimated coefficients on these confounders in our simulation and vulnerability map later, controlling for these ensures that we do not confound the temperature extensive margin adjustment by the other variables. If individuals sort into climate according to income or ethnicity, failing to control for these factors would bias our estimated climate response. Before we move to the discussion of the results, it is important to note the similarities of this approach to that canonical hedonic regression with energy use instead of home values. A recent literature in hedonics discusses the issues and potential solutions to these (e.g., Bajari et al. forthcoming). We will potentially pursue these in future work and here rely on the traditional two-stage approach, due to its computational simplicity and intuitive appeal.

Section 5: Estimation Results

5.1. First Stage Results: The Usage Response to Temperature

As discussed in the previous section, we estimate Equation 1 for each of the 970 ZIP codes that have more than 100 households with at least 47 bills. While we cannot feasibly present the 970 estimated temperature response functions, we can display the temperature response for the highest temperature bin as a map. In Figure 3, we plot the slope of temperature consumption in the highest temperature bin we observe.[8] There are a few points worth noting in this figure. There is tremendous heterogeneity in the temperature response of electricity consumption across ZIP codes, with two sources of variation at work in this picture. First, one would expect that cooler (coastal and mountain) areas of the state would have a flatter response function, as their climates are in the lower parts of the temperature spectrum. If they experienced, say under climate change, a much hotter climate, they might wander up their temperature response curve, parts of which may be currently unobservable to the econometrician, and display a steeper temperature response. Aroonruengsawat and Auffhammer (2012) demonstrate that within a given temperature bin, there is significant variation in temperature response across the state—depending on physical location. They only estimate sixteen distinct temperature response curves, which makes it difficult to examine the source of variation in slope at higher temperatures, which we do in the next stage. To estimate the intensive margin adjustments by the end of century, we follow the approach discussed in AA. Section 6 will show the simulation results with and without incorporating the extensive margin adjustments discussed next.

[8] For completeness, this map plots the temperature response in the highest observed bin for all ZIP codes for which we have any data, including the ZIP codes with fewer than 100 households. We do not use the ZIP codes with fewer than 100 households sampled in any of our formal analysis.

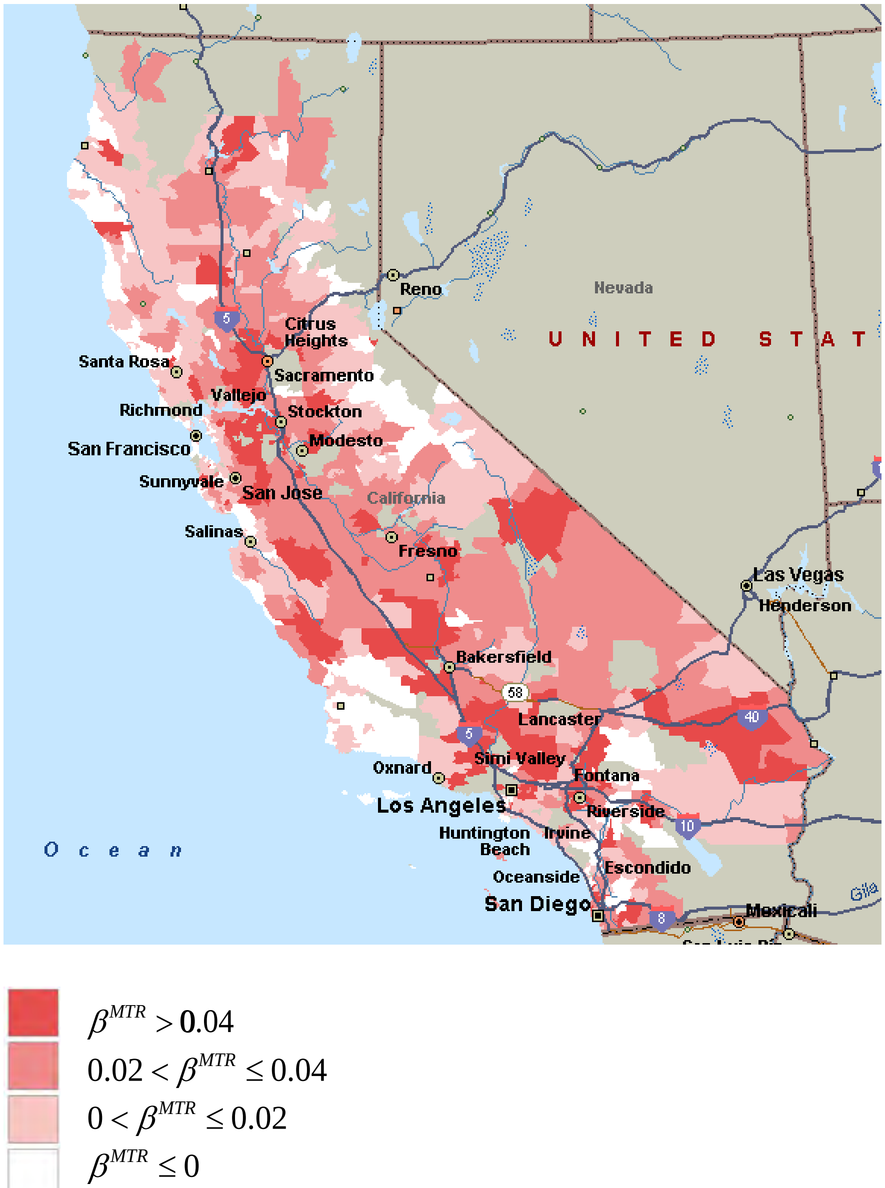

Note: β^{MTR} is the estimated slope for each ZIP code in the highest temperature bin for which data was observed.

Figure 3: Estimated Maximum Temperature Response Coefficients

5.2. Second Stage Results: The Long-Run Response to Temperature

As discussed in Section 4.2, we exploit the 970 estimated temperature response curves and examine whether we can explain variation in temperature response through cross-sectional variation in long-run summer climate and socioeconomics.

The left-hand side variable is our measure of temperature response of electricity consumption multiplied by 100, which we estimated for each ZIP code i in the previous step. On the right-hand side we control for mean summer temperature (June, July, August mean daily temperature) for the years 2003–2006 and observable characteristics of the population at the ZIP code level, as discussed in the data section. Table 3 displays the results from 14 least-squares regressions using the temperature response coefficient for each of the 14 bins and 970 ZIP codes as the left-hand side variable.

It is important to note that individuals do not randomly select into locations. There is evidence from a large literature on economic sorting models that individuals self-sort into locations that match their preferred amenities, given their income level. The regressions in Table 3 therefore carry no causal interpretation whatsoever, but are merely correlations. We would expect that a hotter summer climate (which is correlated with milder winters for most of California) would carry a positive coefficient for the higher temperature bins. While it is not clear that there should be a difference in temperature response in the low temperature bins, one could think that houses in colder areas are slightly better designed for the heating season and more likely to be heated with natural gas instead of baseboard electric heating. This could result in negative coefficients at low temperatures.[9]

The interpretation of the coefficients in Table 3 is not straightforward. The left-hand side variable is itself a regression coefficient which, for values close to zero, approximates the percent change in billing period electricity consumption due to one more day spent in the temperature bin. The coefficient on JJA climate in the regressions in Table 3 therefore indicates that ZIP codes with average June/July/August temperatures from 2003–2006 that are 1°C (1.8°F) higher have temperature response curves that are approximately 0.123 percent steeper for the highest temperature bin. The range of summer temperatures is 21°C (69.8°F), with a standard deviation of 3.5°C (6.3°F).

Warmer summer climate, which for California is consistent with milder winters, leads to a negative coefficient at low temperatures, although the effect is only marginally significant. At higher temperatures however, we observe a strong positive estimate, suggesting that there is a monotonic increase in the slope of the temperature response starting at climate with a summer mean at 71°F (22°C) and higher. We note that there is a small signal in the 59.5°F–63°F (15.3°C–17°C) bin, but it is small in magnitude, and the two adjacent bins do not have a statistically detectable signal

Income per household is significant in all of the models. The variable is measured in $10,000 per household in 2007 dollars. The mean income per household for our data is $48,111, with a standard deviation of $21,222. A one standard deviation change therefore results in a 0.34 percent increase of the slope in the highest temperature bin. A one standard deviation change in income has a 50 percent smaller impact on residential electricity consumption increases than the effect of a one standard deviation increase in mean JJA temperature

[9] We also experimented with using winter temperatures as covariates. Due to collinearity, these were not significant in the models. We also experimented with nonlinear no-parametric functions in temperature, and these are approximately linear, suggesting that our temperature specification is appropriate.

(discussed in the previous paragraph). It would be very useful to determine how important the impact of air conditioner prices and the price of electricity on electricity use on the adoption of air conditioners is, but unfortunately we lack data to conduct such an exercise.

Table 3: Second Stage Regressions of Temperature Response Coefficients by Temperature Bin

	(1)	(2)	(3)	(4)	(5)
Bin	<28.5	28.5–38.5	38.5–43	43–49	49–52.5
JJA Temp	-0.102*	-0.103***	-0.0214	0.00263	-0.00228
	(0.0564)	(0.0272)	(0.0132)	(0.00381)	(0.00305)
Constant	2.210**	3.726***	0.950***	0.147	0.370***
	(0.962)	(0.780)	(0.363)	(0.107)	(0.0705)
Observations	53	449	670	951	970
R-squared	0.155	0.072	0.033	0.007	0.034
Population	Yes	Yes	Yes	Yes	Yes
Income	Yes	Yes	Yes	Yes	Yes
Elevation	Yes	Yes	Yes	Yes	Yes

Model	(6)	(7)	(8)	(9)	(10)
Bin	52.5–56	56–59.5	59.5–63	63–66.5	66.5–71
JJA Temp	2.20e-05	-3.38e-10	-0.0101***	-0.00185	0.00237
	(0.00325)	(7.99e-10)	(0.00329)	(0.00329)	(0.00402)
Constant	0.0796	2.51e-08	0.185**	0.0176	0.233**
	(0.0748)	(2.08e-08)	(0.0735)	(0.0757)	(0.0969)
Observations	970	970	970	970	964
R-squared	0.005	0.009	0.042	0.010	0.052
Population	Yes	Yes	Yes	Yes	Yes
Income	Yes	Yes	Yes	Yes	Yes
Elevation	Yes	Yes	Yes	Yes	Yes

Model	(11)	(12)	(13)	(14)
Bin	71–77	77–82	82–91.5	>91.5
JJA Temp	0.0413***	0.101***	0.151***	0.123**
	(0.00588)	(0.00932)	(0.0219)	(0.0611)
Constant	-0.430***	-1.302***	-2.173***	-0.327
	(0.141)	(0.244)	(0.589)	(1.895)
Observations	964	953	762	350
R-squared	0.074	0.130	0.105	0.055
Population	Yes	Yes	Yes	Yes
Income	Yes	Yes	Yes	Yes
Elevation	Yes	Yes	Yes	Yes

Note: Robust standard errors in parentheses *** p<0.01, ** p<0.05, * p<0.1.

Using the results from Table 3, we calculate the predicted change in slope across temperature bins using the coefficients on summer climate for each bin. For the Figure 4 we use the NCAR PCM1 scenario A2 with analog downscaling, as discussed in AA. We use the predicted change in JJA temperature for 2080–2099 over 2003–2006 from the model and multiply it by the coefficient on climate in Table 3 for each bin. Figure 4 plots the population weighted average across ZIP codes of the estimated temperature response curves from Section 5.1 (black solid line) and compares it to the counterfactual response curves using the results from Table (3) (red line). The average increase over the 2003–2006 estimated slope for the three highest bins in Table 3 is 13 percent. These results are economically and statistically significant. The counterfactual response based on Table 3 allows for changes at each temperature bin and displays a flattening out at the low temperature range and a steepening at higher temperatures. While these results are subject to the caveats discussed above, they suggest adaptation to climate change. We now turn to discussing the magnitude of these effects, as well as their spatial distribution.

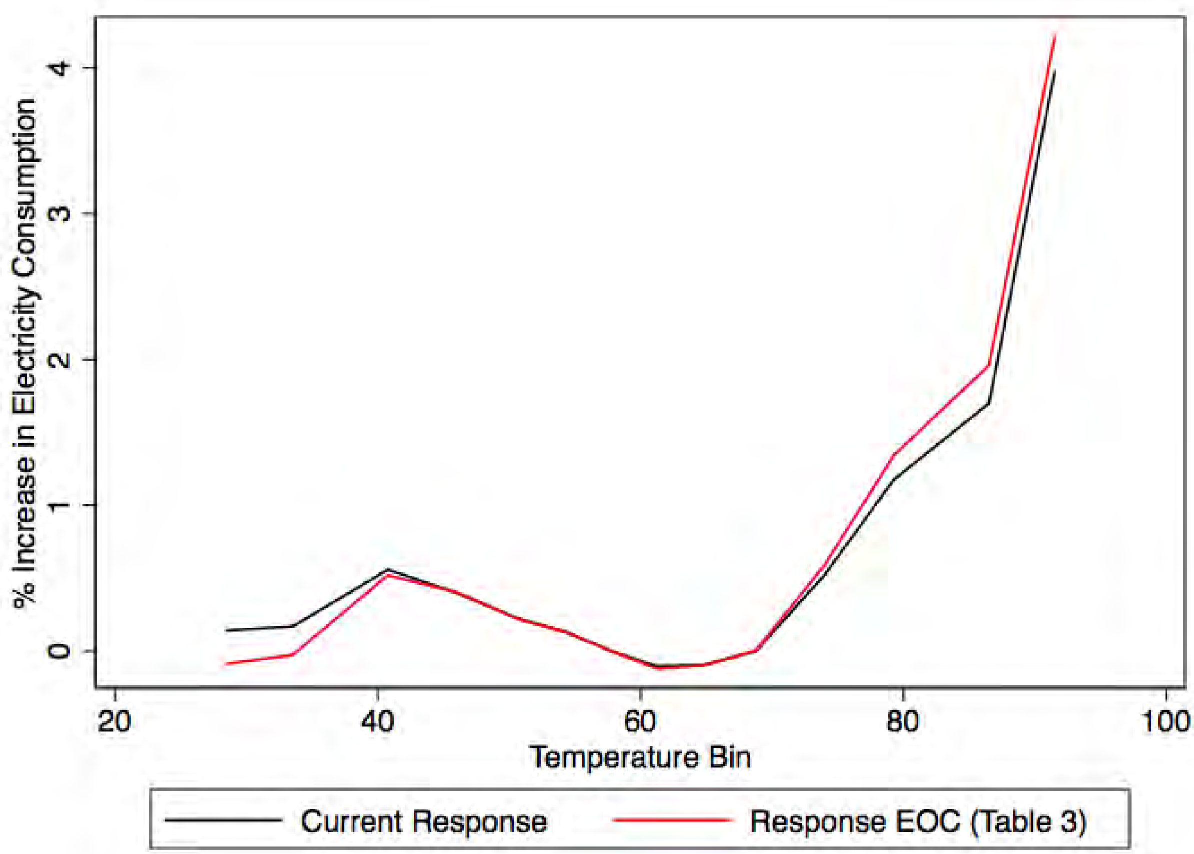

Figure 4: Population-Weighted Temperature Response 2003–2006 versus 2080–2099

Section 6: Electricity Consumption Increases Due to Climate and Population Change

In this section, we simulate the impacts of climate change on electricity consumption under two different emissions scenarios using three different climate models. We conduct three different simulations. The first simulation holds population growth constant and only simulates electricity consumption per household using the first-stage estimates, which does not allow for changes in the extensive margin. In a second simulation we incorporate the extensive margin adjustments conducted in sections 4.2 and 5.2. In a final simulation we allow for population growth. For the first two simulations we provide impacts at the household level, as well as at the aggregate level. The final simulation uses the household-level simulation, allowing for intensive and extensive margin adjustments, and adds population growth—therefore we only provide aggregate forecasts in consumption for this scenario. We can calculate the trajectory of aggregate electricity consumption from the residential sector until the year 2100, which is standard in the climate change literature. Due to space limitations, we only present the result for the 2080–2099 period here.

6.1 Temperature Simulations

The simulation for this section uses the estimated climate response parameters estimated in Section 5.1. Using these estimates as the basis of our simulation has several strong implications. Using the estimated first stage parameters implies that the climate responsiveness of demand within climate zones remains constant throughout the century.

As is standard in this literature, the counterfactual climate is generated by a general circulation model (GCM). These numerical simulation models generate predictions of past and future climate under different scenarios of atmospheric greenhouse gas (GHG) concentrations. The quantitative projections of global climate change conducted under the auspices of the IPCC and applied in this study are driven by modeled simulations of two sets of projections of twenty-first century social and economic development around the world, the so-called “A2” and “B1” storylines in the 2000 *Special Report on Emissions Scenarios* (SRES) (IPCC 2000). The SRES study was conducted as part of the IPCC’s Third Assessment Round, released in 2001.

The A2 and B1 storylines and their quantitative representations represent two quite different possible trajectories for the world economy, society, and energy system, and they imply divergent future anthropogenic emissions, with projected emissions in the A2 being substantially higher. The A2 scenario represents a “differentiated world,” with respect to demographics, economic growth, resource use, energy systems, and cultural factors, resulting in continued growth in global carbon dioxide (CO_2) emissions, which reach nearly 30 gigatons of carbon (GtC) annually in the marker scenario by 2100. The B1 scenario can be characterized as a “global sustainability” scenario. Worldwide, environmental protection and quality and human development emerge as key priorities, and there is an increase in international cooperation to address them as well as convergence in other dimensions. A demographic transition results in global population, peaking around mid-century and declining thereafter, reaching roughly 7 billion by 2100. Economic growth rates are higher than those in A2, so that global economic output in 2100 is approximately one-third greater. In the B1 marker scenario, annual emissions reach about 12 GtC in 2040 and decline to about 4 GtC in 2100.

We simulate demand for each scenario using the NCAR Parallel Climate Model 1 (PCM), the Geophysical Fluid Dynamics Laboratory 2.1. Climate Model (GFDL), and the Centre National de Recherches Météorologiques Climate Model (CNRM) v3 retrieved from the archived statistical downscaling provided by Maurer and Das, archived at the University of California, San Diego. These models were provided to us in their downscaled version for California using the Constructed Analogues algorithms (Maurer and Hidalgo 2008).

To obtain estimates for a percent increase in electricity consumption for the representative household in ZIP code *j* and period *t*+*h,* we use the following relation:

$$\frac{q_{j,t+h}}{q_{j,t}} = \frac{\exp\left(\sum_{p=1}^{k} \hat{\beta}_{pj} D_{pj,t+h}\right)}{\exp\left(\sum_{p=1}^{k} \hat{\beta}_{pj} D_{pj,t}\right)} \tag{3}$$

Figure 5 displays the projected increases in household residential electricity consumption at the ZIP code level for the 970 ZIP codes in our sample for the intensive margin adjustment only. Consistent with AA, this map displays that the ZIP codes in the Central Valley and Southeastern California are projected to experience the largest increases in household electricity consumption.

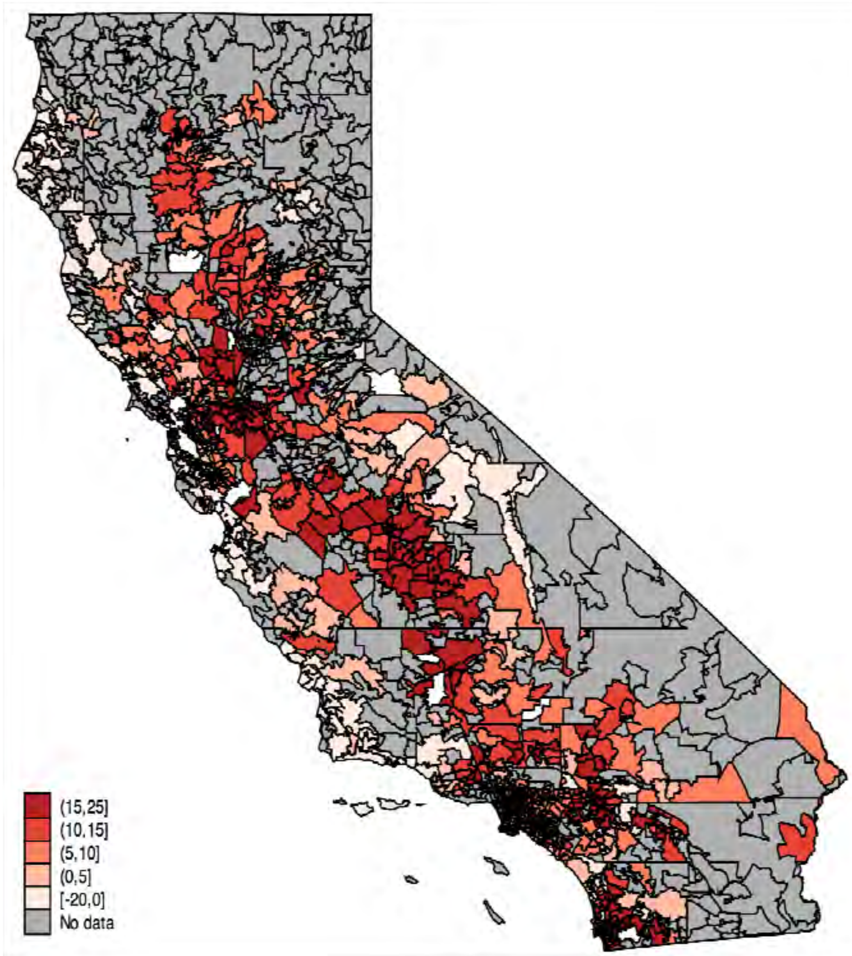

Figure 5: Intensive Margin Adjustment: Projected Percent Increases in Household Electricity Consumption 2080–2099 over 1961–1990 Average Consumption due to a Change in Temperature Increases from GFDLv3 Forced by the SRES A2 Scenario and Downscaled by the Constructed Analogues Algorithm

Even though we do not display the projected increases from the other climate models and remaining periods, it should be noted that the distribution of impacts is very similar to that shown in Figure 5. Figure 6 below uses the estimated extensive margin adjustments as displayed in Table 3 to predict increases in electricity consumption due to a changed temperature response function for each ZIP code. Conceptually this is using the ZIP code equivalent of the red line in Figure 5 to calculate projected increases in electricity consumption. The "darkening" of the central valley and areas in Southern California indicates that the projected impacts are larger, with all ZIP codes experiencing increases in electricity consumption due to this incorporation of the extensive margin adjustment.

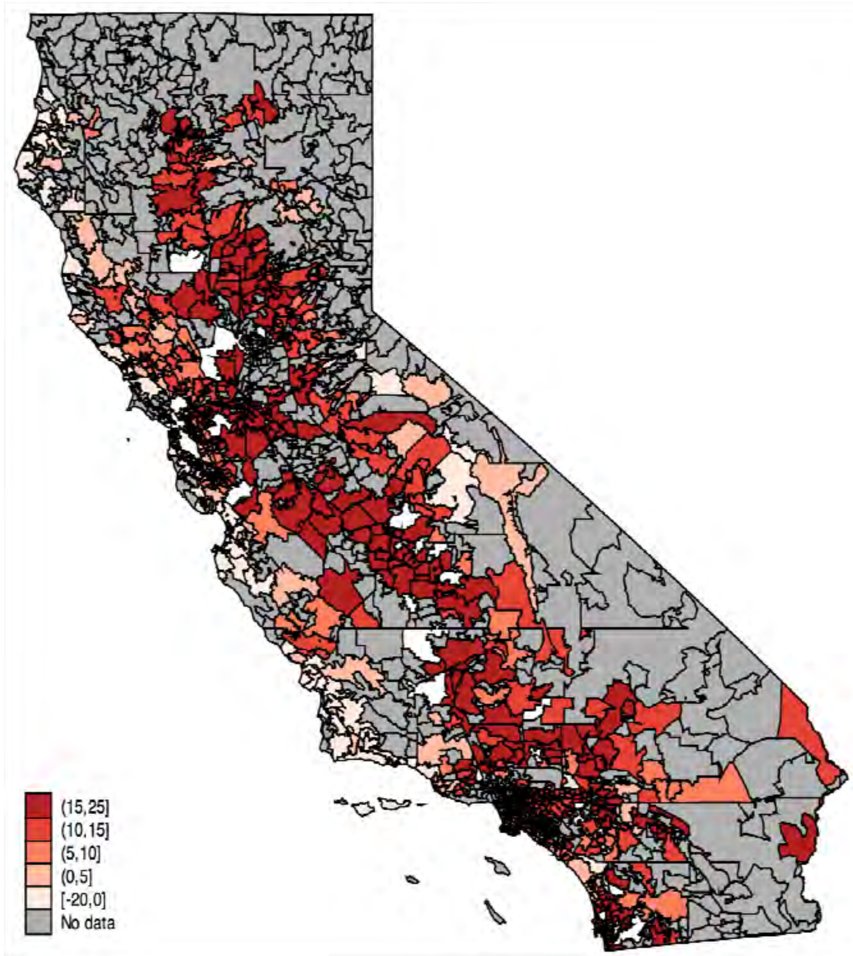

Figure 6: Intensive Plus Extensive Margin Adjustment: Projected Percent Increases in Household Electricity Consumption 2080–2099 over 1961–1990 Average Consumption due to a Change in Temperature Increases from GFDLv3 Forced by the SRES A2 Scenario and Downscaled by the Constructed Analogues Algorithm

While the two maps (figures 5 and 6) are instructive, it is hard to determine how big the overall impact of allowing for extensive margin adjustment is. Table 4 therefore shows the overall population-weighted increases in total electricity consumption predicted by the three climate models and the two SRES scenarios—with and without extensive margin adjustments.

Table 4: Projected Percent Increase in Residential Electricity Consumption due to Climate Change over 1961–1990 Consumption by Climate Model, Forcing Scenario, and Extensive Margin Adjustment Modeling

Climate Model	CNRM	CNRM	GFDL	GFDL	NCAR	NCAR
SRES Scenario	B1	B1	B1	B1	B1	B1
Extensive Margin	No	Yes	No	Yes	No	Yes
2000–19	0.6%	0.8%	1.0%	1.3%	0.2%	0.2%
2020–39	1.2%	1.5%	1.8%	2.2%	0.7%	0.8%
2040–59	1.5%	1.8%	2.2%	2.7%	0.9%	1.0%
2060–79	2.2%	2.7%	2.8%	3.4%	1.2%	1.5%
2080–99	2.6%	3.2%	3.6%	4.3%	1.3%	1.7%

Climate Model	CNRM	CNRM	GFDL	GFDL	NCAR	NCAR
SRES Scenario	A2	A2	A2	A2	A2	A2
Extensive Margin	No	Yes	No	Yes	No	Yes
2000–19	0.69%	1.12%	0.81%	1.43%	0.69%	0.81%
2020–39	1.01%	1.67%	2.05%	3.05%	0.56%	0.72%
2040–59	2.14%	3.20%	3.04%	4.60%	1.37%	1.69%
2060–79	4.08%	6.13%	4.70%	7.03%	2.02%	2.50%
2080–99	6.52%	9.67%	7.74%	11.32%	3.50%	4.25%

Table 4 shows that for SRES scenario A2, if we do not account for extensive margin adjustments, the projected increases in electricity consumption range from 3.50–7.74 percent, depending on the climate model employed. The GFDL model predicts significantly higher increases, as it is a higher sensitivity model. If we allow for the extensive margin adjustment, the increases range from 4.25 percent to 11.32 percent. For the GFDL model the overall increase increases by more than 3.5 percentage points, which is a significant increase. It is not surprising that the higher-sensitivity models (GFDL and CNRM) show higher increases along the extensive margin adjustment dimension as well, as these are driven also by the higher projected temperature increases. For the lower-sensitivity NCAR model, the difference is only 0.75 percentage points by the end of the century—yet this represents a 21.42 percent proportional increase over the intensive margin-only adjustment.

6.2 Temperature and Population Simulations

California has experienced an almost seven-fold increase in its population since 1929 (BEA 2008), and California's population growth rate over that period (2.45 percent) was more than double that of the national average (1.17 percent). Over the past 50 years California's population has grown by 22 million people to almost 37 million in 2007 (BEA 2008). To predict what the trajectory of California's population will look like until the year 2100, many factors have to be taken into account. The four key components driving future population are net international migration, net domestic migration, mortality rates, and fertility rates. The State of California provides forecasts 55 years into the future, which is problematic, since we are interested in simulating end-of-century electricity consumption. The Public Policy Institute of California has generated a set of population projections until 2100 at the county level, and we obtained these from Sanstad et al. (2009).

The three sets of projections developed for California and its counties are designed to provide a subjective assessment of the uncertainty of the state's future population. The projections present three very different demographic futures. In the low series, population growth slows as birth rates decline, migration out of the state accelerates, and mortality rates show little improvement. In the high series, population growth accelerates as birth rates increase, migration increases, and mortality declines. The middle series, consistent with (but not identical to) the California Department of Finance projections, assumes future growth in California will be similar to patterns observed over the state's recent history—patterns that include a moderation of previous growth rates but still large absolute changes in the state's population. In the middle series, international migration flows to California remain strong to mid-century and then subside, net domestic migration remains negative but of small magnitude, fertility levels (as measured by total fertility rates) decline slightly, and age-specific mortality rates continue to improve.

The high projection is equivalent to an overall growth rate of 1.47 percent per year and results in a quadrupling of population to 148 million by the end of the century. The middle series results in a 0.88 percent annual growth rate and 2.3-fold increase in total population. The low series is equivalent to a 0.18 percent growth rate and results in a population 18 percent higher than today's. Projections are available at the county level and not at the ZIP code level. We therefore assume that each ZIP code in the same county experiences an identical growth rate. We only use the medium population growth scenario in this paper.

Table 5 displays the simulated aggregate electricity demand given the medium population growth scenarios under climate change and the full intensive and extensive margin adjustment. Consistent with AA, it is not surprising to see that population growth has much larger consequences for simulated total electricity consumption compared to climate uncertainty or price uncertainty. The simulations for the low forcing scenario B1 and the medium population growth scenario show a 134–141 percent increase in consumption, which is largely due to projected increases in population. For the A2 scenario, the predicted increases range from 141–160 percent increases over the 1960–1990 baseline weather year 2000 population. This, unsurprisingly, stresses that population trajectories are much bigger drivers of residential electricity demand than climate change.

Table 5: Simulated Percent Increase in Residential Electricity Consumption Relative to 1961–1990 for the Middle Population Scenario and Extensive Plus Intensive Margin Adjustment of the Temperature Response Function

Climate Model	CNRM	GFDL	NCAR	CNRM	GFDL	NCAR
SRES Scenario	A2	A2	A2	B1	B1	B1
Extensive Margin	Yes	Yes	Yes	Yes	Yes	Yes
2000–19	14%	14%	13%	13%	14%	13%
2020–39	43%	45%	42%	43%	44%	42%
2040–59	76%	79%	73%	73%	75%	72%
2060–79	112%	115%	104%	105%	106%	102%
2080–99	154%	160%	141%	138%	141%	134%

6.3 Heterogeneity in Temperature Response and Impacts

Vulnerability is commonly defined as the interaction between exposure, sensitivity, and adaptive capacity. The maps displayed in Section 6.1 combine exposure to climate change–driven temperature increases with each ZIP code's specific sensitivity to climate change. The projections, which allow for intensive and extensive margin adjustment, simulate how much more households will operate and in some areas engage in the operation of new capital equipment in the form of air conditioners. These maps could therefore be interpreted as maps of vulnerability of households to climate change–induced changes in electricity consumption at the ZIP code level. However, these maps fail to provide insight into two crucial questions.

First, what these maps do not provide estimates of is the true adaptive capacity of households. In other words, is adaptation to climate change along the intensive and extensive margins income dependent? One direct test of this is to run the cross-sectional models in tables 3 and 4 with an interaction term between climate and income. In Table 6 we add an interaction term between income and climate into a regression where we use temperature response in the highest temperature bin for a given ZIP code as the dependent variable. The interaction between income and climate is highly significant. The coefficient on climate is not statistically significant by itself, but jointly significant with the interaction term at 1 percent level. This regression provides some suggestive evidence that residents in ZIP codes of more wealthy areas have a higher response to hot temperatures (e.g., a higher penetration of air conditioners) than do residents in poorer, equally hot, ZIP codes. This suggests that adaptive capacity in this sector is likely tied to income. While our results are not causal, this is consistent with economic theory.

Table 6: Interaction between Climate and Income and Maximum Temperature Response

JJA Temp	0.0532
	(0.0818)
Elevation	-0.000408***
	(9.74e-05)
Household Income	-1.063***
	(0.400)
JJA Temp * HH Income	0.0589***
	(0.0202)
Constant	0.178
	(1.672)
Observations	970
R-squared	0.154

Note: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Second, it is instructive to break down the projected climate-driven percent increases in electricity consumption by different characteristics of the population. Here we pick out four specific dimensions: Income (Figure 7), %Caucasian (Figure 8), %Hispanic (Figure 9) and %African American (Figure 10). While these scatterplots are again not causal, they are informative. There is a statistically significant negative correlation between income and projected increases in electricity consumption for the 2080–2099 period. This indicates that after accounting for intensive and extensive margin adjustments, wealthier ZIP codes will have experienced slightly lower increases in residential electricity consumption compared to their less wealthy counterparts. The coefficient on the regression line depicted in the figure is consistent with 0.44 percent lower increases in household electricity consumption for each $10,000 in income.

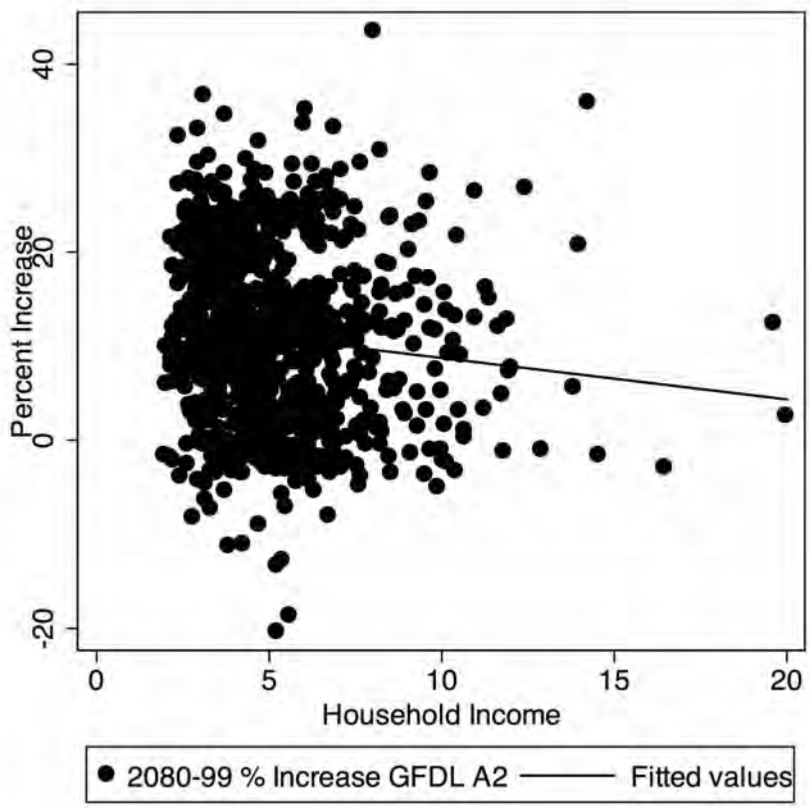

Figure 7: Projected Percent Increases in Household Electricity Consumption 2080–2099 over 1961–1990 Average Consumption Due to Change in Temperature Increases from GFDLv3 Forced by the SRES A2 Scenario and Downscaled by the Constructed Analogues Algorithm Against Average ZIP Code Income

Figure 8 displays again a slightly negative, yet statistically significant, correlation between predicted increases and the share of the population that is Caucasian. Each 10 percent higher share in this fraction is consistent with a 0.38 percent lower increase in electricity consumption. The slightly negative correlation coefficient between each ZIP code's share of African American population depicted in Figure 9 is not statistically different from zero.

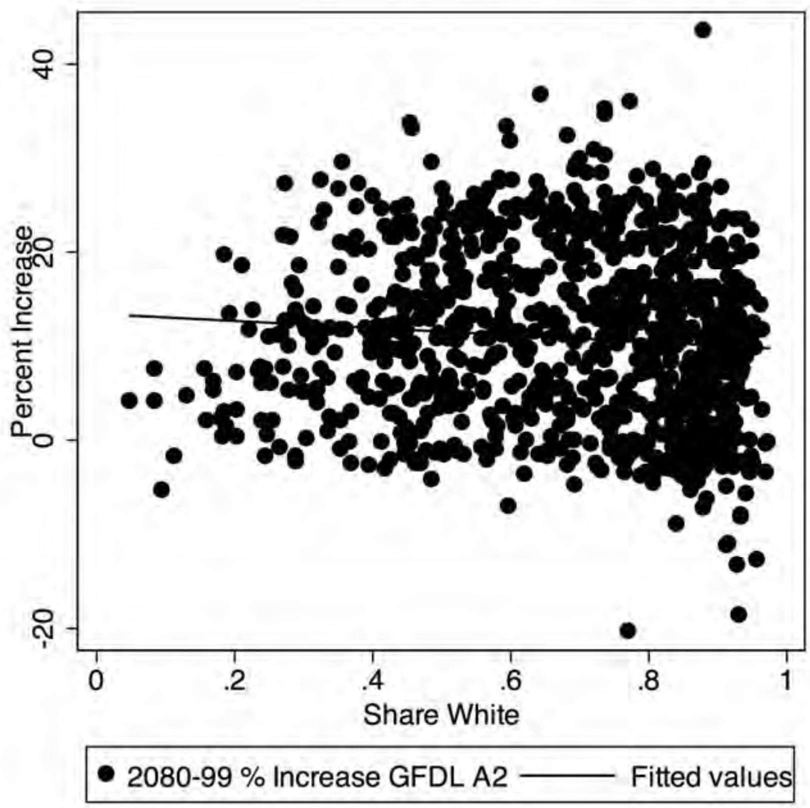

Figure 8: Projected Percent Increases in Household Electricity Consumption 2080–2099 over 1961–1990 Average Consumption Due to a Change in Temperature Increases from GFDLv3 Forced by the SRES A2 Scenario and Downscaled by the Constructed Analogues Algorithm Against Zip Code Share of Caucasian Population

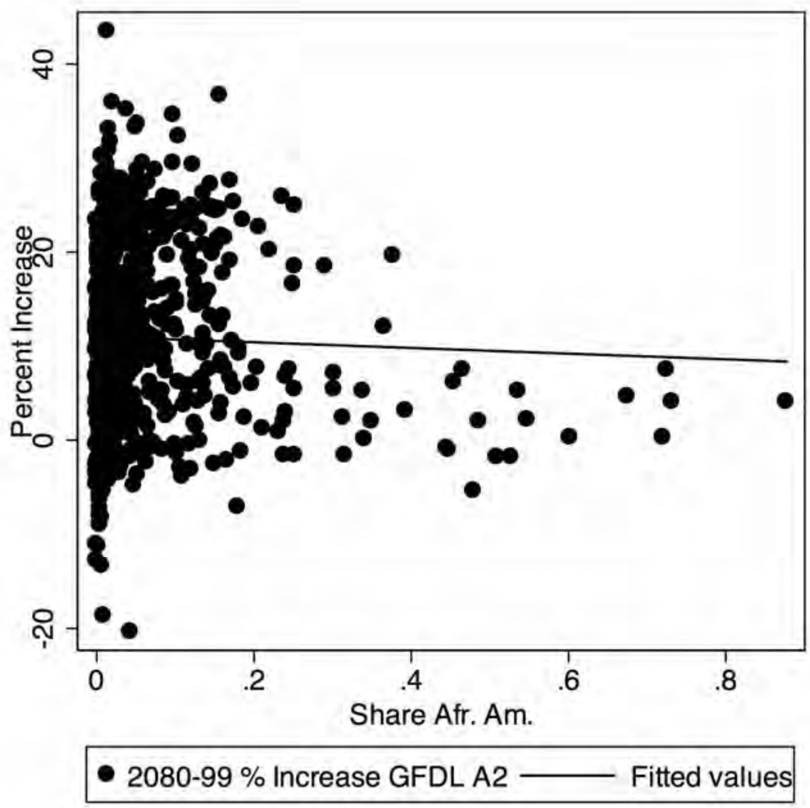

Figure 9: Projected Percent Increases in Household Electricity Consumption 2080–2099 over 1961–1990 Average Consumption Due to a Change in Temperature Increases from GFDLv3 Forced by the SRES A2 Scenario and Downscaled by the Constructed Analogues Algorithm Against ZIP-Code Share of African American Population

Figure 10 displays the final raw correlation between the share of Latino population and predicted increases in electricity consumption. There is a statistically significant and sizable correlation between these two variables. The regression suggests that for each 10 percent increase in the share of Latino population, the predicted increases in electricity consumption increases by 1 percent.

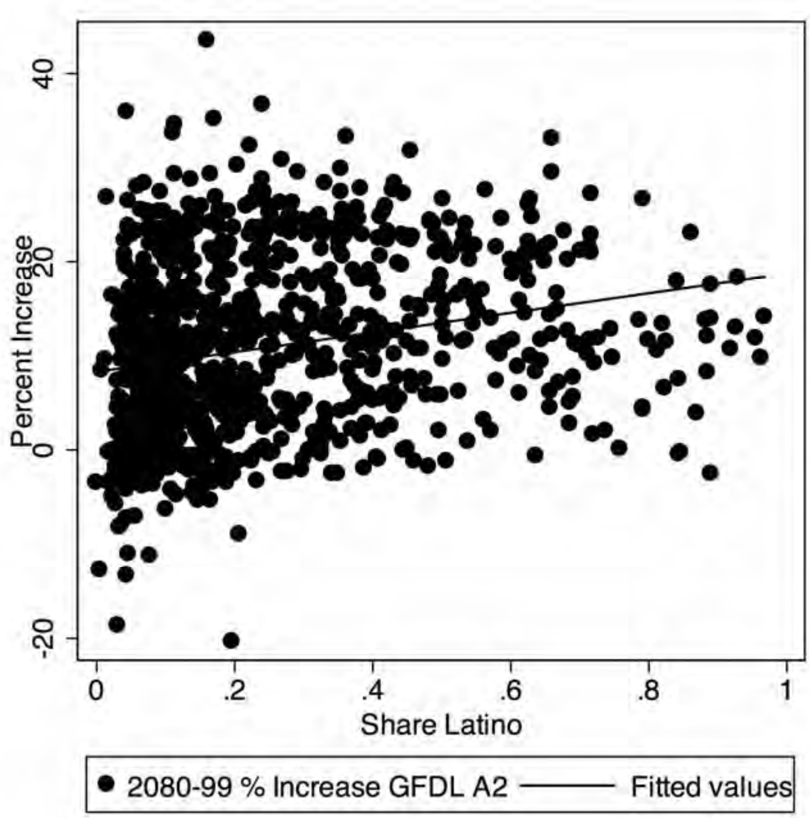

Figure 10: Projected Percent Increases in Household Electricity Consumption 2080–2099 over 1961–1990 Average Consumption Due to a Change in Temperature Increases from GFDLv3 Forced by the SRES A2 Scenario and Downscaled by the Constructed Analogues Algorithm Against ZIP code Share of Latino Population

While these four scatter plots and corresponding estimated regression lines are not causal, if we assume that the distribution of income and ethnic groups stays constant, the results suggest that ZIP codes that contain more Caucasian and wealthy people are likely to experience lower increases in their climate-driven electricity consumption. The opposite is true for ZIP codes with larger shares of low-income households and Latino population. Much of this is due to the spatial distribution of the population. If individuals move in response to climate change, these correlations break down. They are suggestive nonetheless.

Section 7: Conclusions

In the residential sector, one of the most obvious modes of adaptation to higher temperatures due to climate change is the increased demand for cooling and decreased demand for heating in the built environment. Due to its mild climate and heavy reliance on natural gas, California's residential sector uses relatively little electricity for heating. It is therefore expected that the demand for electricity will increase as households operate existing air conditioners more frequently, and in many regions will install air conditioners where there currently are few. This paper provides reduced form estimates of changes in electricity consumption due to increased use of installed cooling equipment under a hotter climate. This study adds to the literature by incorporating the change in temperature responsiveness due to likely increases in air conditioner penetration under climate change using a two-stage method. The paper shows that accounting for these capital investments will lead to statistically and economically significantly higher projections of electricity consumption. These increases in projected electricity consumption were mapped to the ZIP codes in the study data, which can be interpreted as a map of vulnerability of the residential sector to climate-driven increases in electricity demand. The paper shows suggestive evidence that more Caucasian and wealthy ZIP codes are projected to experience relatively smaller increases in consumption, while ZIP codes with a higher share of Latino population and less-wealthy households are projected to experience larger increases in consumption.

It is important to note that we do not and cannot model changes in electricity consumption due to improvements in the efficiency of heating and cooling equipment and/or buildings. These effects will be offsetting some of the gains in consumption outlined in this paper. Further, the extensive margin adjustments in this paper cannot meaningfully control for changes in urban form, urban heat island effects, or other variables potentially leading to a higher response, which may be correlated with temperature. We leave the study of these effects to future work, but caution the reader that the second-stage estimates are correlations and should not be given a causal interpretation.

References

Aroonruengsawat, Anin, and Maximilian Auffhammer. 2012. Correction: Impacts of Climate Change on Residential Electricity Consumption: Evidence From Billing Data. California Energy Commission White Paper.

Bajari, Patrick L., Jane Corinne Cooley, Kyoo-il Kim, and Christopher Timmins. Forthcoming. A Rational Expectations Approach to Hedonic Price Regressions with Time-Varying Unobserved Product Attributes: The Price of Pollution. *American Economic Review*.

Baxter, L. W., and K. Calandri. 1992. "Global warming and electricity demand: A study of California." *Energy Policy* 20(3): 233–244.

Biddle, Jeff. 2008. "Explaining the spread of residential air conditioning, 1955–1980." *Explorations in Economic History* 45: 402–423.

BEA (Bureau of Economic Analysis). 2008. *Regional Economic Accounts*. Washington, D.C. www.bea.gov/regional/spi/default.cfm?series=summary.

Cline, W. R. 1992. *The economics of global warming*. Washington: Institute for International Economics.

Crowley, C., and F. Joutz. 2003. Hourly electricity loads: Temperature elasticities and climate change. In *23rd U.S. Association of Energy Economics North American Conference*.

Deschênes, O., and M. Greenstone. 2011. "Climate Change, Mortality and Adaptation: Evidence from annual fluctuations in weather in the U.S." *American Economic Journal: Applied Economics. 3: 152–185.*

Energy Information Administration (EIA). 2008. State Energy Data System. Washington, D.C. www.eia.doe.gov/emeu/states/_seds.html.

Energy Information Administration. 2011. "Air conditioning in nearly 100 million U.S. homes." http://www.eia.gov/consumption/residential/reports/air_conditioning09.cfm?src=email.

Franco, G., and A. Sanstad. 2008. "Climate change and electricity demand in California." *Climatic Change* 87:139–151.

Hayhoe, K., M. Robson, J. Rogula, M. Auffhammer, N. L. Miller, J. VanDorn, and D. Wuebbles. 2010. "An integrated framework for quantifying and valuing climate change impacts on urban energy and infrastructure: A Chicago case study." *Journal of Great Lakes Research*. 36: 94–105.

Hanemann, W. 1984. "Discrete/Continuous Models of Consumer Demand." *Econometrica* 52: 541–561.

Houthakker, H. S., and L. D. Taylor. 1970. *Consumer Demand in the United States: Analyses and Projections*. Cambridge, Mass: Harvard University Press.

Intergovernmental Panel on Climate Change (IPCC). 2000. *Emissions Scenarios*. Cambridge, UK: Cambridge University Press.

Linder, P., M. Gibbs, and M. Inglis. 1987. *Potential impacts of climate change on electric utilities.* Report 88-2. Albany, New York: New York State Energy Research and Development Authority.

Mansur, E., R. Mendelsohn, and W. Morrison. 2008. "Climate change adaptation: A study of fuel choice and consumption in the U.S. energy sector." *Journal of Environmental Economics and Management* 55(2): 175–193.

Maurer, E. P., and H. G. Hidalgo. 2008. "Utility of daily vs. monthly large-scale climate data: An intercomparison of two statistical downscaling methods." *Hydrology and Earth System Sciences* 12:551–563.

McNeil, M., and V. Letschert. 2008. Future Air Conditioning Energy Consumption in Developing Countries and what can be done about it: The Potential of Efficiency in the Residential Sector. LBNL Working Paper. http://escholarship.org/uc/item/64f9r6wr.

McNeil, M., and V. Letschert. 2010. "Modeling diffusion of electrical appliances in the residential sector." *Energy and Buildings* 42: 783–790.

Mendelsohn, R. 2003. The Impact of Climate Change on Energy Expenditures in California. Appendix XI in Wilson, Tom, and Larry Williams, Joel Smith, Robert Mendelsohn. *Global Climate Change and California: Potential Implications for Ecosystems, Health, and the Economy.* Consultant report 500-03-058CF to the Public Interest Energy Research Program, California Energy Commission.

Miller, N. L., K. Hayhoe, J. Jin, and M. Auffhammer. 2008. "Climate change, extreme heat, and electricity demand." J. *Applied Meteorology and Climatology* 47(6): 1834–1844.

Rapson. D. 2011. "Durable Goods and Long-Run Electricity Demand: A Case Study of Air Conditioner Purchase Behavior." Working Paper, UC Davis.

Reiss, P. C., and M. W. White. 2005. "Household Electricity Demand Revisited." *The Review of Economic Studies*. 72:853–883.

Rosenthal, D., H. Gruenspecht, and E. Moran. 1995. "Effects of global warming on energy use for space heating and cooling in the United States." *The Energy Journal* 16:77–96.

Sailor, D., and A. Pavlova. 2003. "Air conditioning market saturation and long-term response of residential cooling energy demand to climate change." *Energy* 28 (9): 941–51.

Sanstad, A., H. Johnson, N. Goldstein, and G. Franco. 2009. *Long-Run Socioeconomic and Demographic Scenarios for California.* PIER Report CEC-500-2009-013-D.

Schlenker, W., and M. Roberts. 2006. "Nonlinear Effects of Weather on Corn Yields." *Review of Agricultural Economics* 28(3): 391–398.

U.S. Environmental Protection Agency. 1989. *The potential effects of global climate change on the United States.* Joel B. Smith and Dennis Tirpak, eds. Washington: Environmental Protection Agency.

Glossary

AA	Aroonruengsawat and Auffhammer
AC	air conditioner
BEA	Bureau of Economic Analysis
CDD	cooling degree days
CEC	California Energy Commission
CNRM	Centre National de Recherches Météorologiques Climate Model
CO_2	carbon dioxide
EIA	Energy Information Administration
GCM	general circulation models
GFDL	Geophysical Fluid Dynamics Laboratory
GHG	greenhouse gas
GIS	geographical information system
GtC	gigatons of carbon
HH	household
IPCC	Intergovernmental Panel on Climate Change
JJA	June, July, August
kWh	kilowatt-hour
NCAR	National Center for Atmospheric Research
NCDC	National Climate Data Center
NOAA	National Oceanic and Atmospheric Administration
PCM1	parallel climate model
SEDS	State Energy Data System
SRES	Special Report on Emissions Scenarios
UCSD	University of California, San Diego