



Energy Institute WP 299

Electric Heating and the Effects of Temperature on Household Electricity Consumption in South Africa

Susanna Berkouwer

February 2019

Energy Institute at Haas working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to review by any editorial board. The Energy Institute acknowledges the generous support it has received from the organizations and individuals listed at <http://ei.haas.berkeley.edu/support/>.

© 2019 by Susanna Berkouwer. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit is given to the source.

Electric Heating and the Effects of Temperature on Household Electricity Consumption in South Africa

Susanna Berkouwer *

February 5, 2019

Abstract

Household energy consumption in developing countries is expected to surge in the coming decades. Yet little is known about how temperature drives household energy demand in developing countries. This paper uses 132,375,282 hourly electricity consumption observations from 5,975 households in South Africa to estimate the effects of temperature on household electricity consumption. The estimates flexibly identify a constant log-linear temperature response: for every 1°C increase in temperature, electricity consumption decreases by 4.1% among temperatures below the heating threshold but increases by 12.2% among temperatures above the cooling threshold. This relationship is driven more strongly by seasonal than hourly temperature changes. Holding all else constant, a 3.25°C increase in temperatures would reduce electricity consumption by 1,093.4 kWh (6.2%) per year per household. Widespread use of electric heating due to limited residential gas heating infrastructure likely drives this. These results point to important regional heterogeneity in how temperature increases may affect household energy demand in the coming decades.

JEL codes: 013, Q4, Q54, C55, C33

Keywords: Economic Development and Energy, Electricity, Energy, Climate, Large Data Sets: Modeling and Analysis, Panel Data

*Department of Agricultural and Resource Economics, University of California, Berkeley; Energy Institute at Haas. 247-A University Hall, Berkeley, CA 94720. E-mail: berkouwer@berkeley.edu. I am grateful to Meredith Fowlie for valuable guidance and discussions. I thank Kelsey Jack, Grant Smith, and Eng. Raj Chetty for generously helping to make this project possible. I also thank Maximilian Auffhammer, Patrick Baylis, Joshua Blonz, Fiona Burlig, Tamma Carleton, Richard Chandler, Sergio Correia, Lucas Davis, Walter Graf, Jay Taneja, Matthew Woerman, Catherine Wolfram, and two anonymous referees for helpful discussions and comments.

1 Introduction

In the next several decades, energy use in emerging markets is expected to rise dramatically. From 2007 to 2035, energy consumption is predicted to grow by 84% in non-OECD countries compared to 14% in OECD countries (Wolfram et al. (2012)). However, despite these large predicted increases, little is known about how the electricity consumption profiles of middle and higher income households in developing countries respond to temperature changes. In addition, while much research in this area has focused on the marginal increase in energy demand caused by an increase in air-conditioning usage when temperatures reach higher extremes, less is known about the response to temperature increases at lower temperatures.

Household temperature responses in emerging markets more commonly include electric heating rather than gas heating, suggesting that energy consumption responses at lower temperatures may be substantially different in these regions than in Europe and North America. In many OECD countries, gas constitutes the primary source of space heating. In the U.S., natural gas (49%) and fuel oil, kerosene, and propane (10%) constitute the bulk of primary space heating for households, with electricity the primary heating source for only 38% of households, according to the EIA Residential Energy Consumption Survey (2015). Dunbabin et al. (2015) find that the primary fuel source of heating in the UK is also gas, and only 10% of British households supplement with electrical heating in cold winter months. While data are limited, these figures may look different in developing countries, where residential gas infrastructure is not as prevalent. According to the World Bank (2011), as incomes rise, the poorest households acquire access to electricity well before gas heating. In South Africa, only 3% of households use natural gas as their primary source of heating, while 67% cite electricity as their primary source of space heating, according to the 2011 South African Census. Electricity is also the primary source of heating in Brazil (Coelho et al. (2014)). Poor insulation and unreliable infrastructure frequently found in lower-income countries can lead to higher levels of energy consumption at low temperatures than in higher-income countries. This suggests that we may see significantly higher energy consumption

at lower temperatures. Gupta (2014) uses daily panel data at the state level from India to document that household electricity demand has a stronger negative response to higher temperatures in cooler regions, where households adapt to colder temperatures using electric heating equipment. Yet, little is known about how temperature affects short-term household-level energy consumption in settings where a majority of households employ electricity rather than natural gas for space heating.

This paper fills this gap by providing a comprehensive understanding of the effect of temperature on household electricity demand in South Africa. Like in much of Africa, penetration of household air-conditioning units in South Africa is limited, while electric heating is widespread. To the best of my knowledge, this is the first paper that uses high-frequency household-level electricity consumption data from a developing country. This paper uses 132,375,282 hourly household electricity consumption observations from 2010 to 2013 for 5,975 households in Johannesburg to study the impact of temperature on electricity consumption. The analysis identifies two main results.

First, households use substantially more electricity at cold temperatures than at medium temperatures. An algorithm designed to search across the temperature range in $0.5^{\circ}C$ intervals to minimize the root mean squared error identifies a heating threshold of $23^{\circ}C$ and a cooling threshold of $30.5^{\circ}C$. Across the entire temperature range $2-23^{\circ}C$, households consistently consume 4.1% less electricity for every $1^{\circ}C$ increase in outdoor temperatures. Consumption remains largely constant between $23-30.5^{\circ}C$, and then appears to increase sharply for temperatures greater than $30.5^{\circ}C$. Aggregating the effects across the distribution of temperatures, the heating response dominates the cooling response. Holding all else constant, including the household rate of ownership of household appliances such as electric heating and air conditioning, an increase in temperatures of $3.25^{\circ}C$ by mid-century would cause a marginal decrease in electricity consumption of 1,093.4 kWh per year per household, or 6.2% relative to baseline levels.

Second, household temperature responses vary strongly by season, and household temperature response becomes much more muted after controlling for seasonal temperature variation.

Particularly during the summer months, household electricity consumption does not appear to respond strongly to temperature variation in the short-term. By 2040, Johannesburg is expected to experience 35 additional days per year where the maximum temperature reaches above $27^{\circ}C$, with an additional significant increase in the number of days where the maximum temperature reaches above $32^{\circ}C$ (Garland et al. (2015)). However, if households do not respond to daily fluctuations as much as they respond to increases in average seasonal temperatures, then current estimates that are based on large fluctuations in daily weather may overestimate the extent to which increases in daily temperature fluctuations will drive an increase in energy consumption.

This paper proceeds as follows. Section 2 provides an overview of the existing literature in this field. Section 3 provides context around the South African household electricity pricing and market structure. Section 4 presents the data and Section 5 presents the identification strategy. Section 6 discusses the results and presents several robustness and heterogeneity checks to verify these results. Section 7 simulates household responses to predicted temperature increases under climate change. Section 8 concludes.

2 Literature review

A recent surge in energy economics research has investigated the predicted growth in household energy consumption in developing countries in the next several decades. Wolfram et al. (2012) predict that energy consumption will grow by 84% in non-OECD countries compared to 14% in OECD countries between 2007 and 2035. Shah et al. (2015) record that air conditioning ownership as a fraction of the number of households in China has increased from less than 5% in 1990 to more than 100% in 2010, and that air conditioning ownership in other emerging markets are growing at 10%-15% annually. A changing climate is likely to affect electricity demand in the coming decades. Exactly how temperatures affect electricity consumption in a development setting, however, remains an understudied research question.

A growing literature documents short-term household energy responsiveness to temperature

changes, but much of this is based on data from the U.S. and Europe. Engle et al. (1986) evaluate various approaches to nonparametric and semiparametric modeling, and apply these to data from utilities based in the U.S. states of Missouri, Georgia, Connecticut, and Washington. More recently, Auffhammer et al. (2016), Aroonruengsawat and Auffhammer (2011), Deschenes and Greenstone (2011), and Crowley and Joutz (2003) study hourly household electricity consumption and find that total electricity consumption will increase as a result of climate change. Much of the literature described the household electricity temperature response function as being U-shaped, with energy consumption being especially high at very low temperatures (when households use electric heating) and at very high temperatures (when households use electric cooling). Wang and Chen (2014) and Petri and Caldeira (2015) both study cooling and heating responsiveness to climate change in the U.S. and find that some cities would experience a net increase in source energy use for cooling and heating by the 2080s while other cities would experience a net reduction in source energy use, due to geographic variation. Wenz et al. (2017) use high-frequency temperature and electricity load data to study the heterogeneity in these temperature response functions across a range of 35 countries, but their data are limited to countries in Europe only.

Studies that look at the relationship between temperature and energy consumption in a development setting often do so using engineering models, monthly data, or using data aggregated at the city or region level. Shah et al. (2015) use engineering models to calculate that air conditioning can account for up to 60% of total load on peak summer days in hot climates such as Delhi, and up to 30% in warm climates such as California. Waite et al. (2017) conduct a comprehensive meta-analysis on electricity consumption across 18 OECD and 17 non-OECD electricity utilities to identify a heating change point and a cooling change point for each city. While they identify a strong cooling relationship in many cities, with electricity consumption increasing sharply at higher temperatures, they find only moderate effects of electric heating on electricity consumption at lower temperatures, and only in a small number of cities. Gupta (2014) uses state-level daily electricity data to document a positive temperature response function at higher temperatures but

a negative response to higher temperatures in cooler regions, attributable to households adapting to colder temperatures using electric heating equipment. However, there is often a focus on the positive response at higher temperature, attributable to increased usage of air conditioning units. Auffhammer (2014) studies the extensive margin of air conditioner ownership and identifies increased household air conditioning use as a significant driver of future increased energy consumption in developing countries. On the intensive margin, Davis and Gertler (2015), use monthly billing data from Mexico and find that electricity consumption increases with temperature given current levels of air conditioning.

Even less evidence is available on electricity consumption in South Africa. Chikobvu and Sigauke (2010) and Chikobvu and Sigauke (2013) study temperature drivers of daily electricity demand to generate daily peak load forecasts, but they use daily maximum hourly demand aggregated at the utility level. Jack and Smith (2016) study electricity consumption among prepaid and postpaid customers in Cape Town, but do so using monthly billing data only. Understanding the relationship between temperature and energy consumption in a development setting in particular is important because economic and geographic differences can cause household incomes, household appliance ownership patterns, the power generation mix, the climate profile, and the energy infrastructure to be substantially different in non-OECD countries, which may all affect this relationship. In addition, technological and institutional hurdles in the electricity sector have historically made it more difficult to access high-frequency electricity and temperature data in development settings, meaning we currently have a less detailed understanding of this relationship than we do in OECD countries. To the best of my knowledge, this is the first paper to use hourly household level data to study the effect of temperature on electricity consumption in a development setting.

3 Background

According to the 2011 South African Census, 59% of South African households use electricity as their primary energy source for heating, with only 2.5% of households primarily using gas,

and 26% using traditional energy sources such as wood, paraffin, coal, or animal dung. 83% of households connected to electricity use electric appliances to heat water. This paper studies 5,975 households that purchase electricity from Eskom, South Africa’s national electric utility. Eskom currently generates approximately 95% of the electricity used in South Africa and operates a vast network of electricity generation, transmission, and distribution operations throughout Africa. A large share of electricity consumers in South Africa, including a majority of Eskom customers in the city of Cape Town, pay monthly installments prior to the consumption of electricity (Jack and Smith (2015)). These customers are generally referred to as *pre-paid* customers. The remaining customers are *post-paid* customers that pay for the electricity that they consumed at the end of each month. All 5,975 households in this dataset are post-paid customers.

Eskom employs a standard block rate tariff scheme¹, where the price per kWh increases by roughly 58% once households exceed a threshold of 600 kWh of electricity consumption during that month. Prices are changed once per year on July 1st by an amount that is previously agreed upon through negotiations between Eskom and *NERSA*, the National Energy Regulator of South Africa². Any discontinuous jump on this date will be absorbed by the monthly fixed effects included in the main specification. Auffhammer (2018) finds that estimates for the coefficients of the effect of temperature on energy demand are identical when price is included in a number of different ways, as well as when it is excluded from the estimating equation altogether, suggesting that price response are negligible in temperature estimation. The specifications in this paper therefore exclude price.

4 Data

I study hourly household electricity consumption by 5,975 households in the province of Gauteng, South Africa from 1 January 2010 to 31 March 2013. All households are located in the relatively

¹While Eskom has considered implementing a time-varying pricing structure to reduce peak demand in the past, and have conducted some small-scale pilot testing, they currently do not employ any time-varying pricing schemes.

²For example, on July 1st 2016, NERSA increased prices by 12.7%

higher-income Fourways and Sandton suburbs roughly 20 kilometers north of downtown Johannesburg. Households in this neighborhood often earn upwards of *R*307,601 (\$37,667) per year according to the 2011 Census, placing many in the top 7.8% of household income earners in South Africa. Due to the higher average income of the households, the extent to which these households experience power outages is minimal, as discussed in Figure A1.

Air conditioner ownership is also likely to be higher among this group, and their electricity consumption for cooling purposes is therefore likely an upper bound. Electricity is the primary energy for space heating for 75% of households in Gauteng Province, which includes Johannesburg. Some households enter (leave) the data later (earlier) than the full sample period. Figure A3 therefore estimates the main specification using only data from 4,651 households between 1 January 2011 to 31 December 2012, for which the data constitute a balanced panel. The results are consistent.

4.1 Electricity data

Household electricity consumption is recorded electronically using identical pieces of smart metering technology installed in each household. For each household I remove all observations prior to the first hour where they consume a non-zero amount of electricity. These earlier data likely reflect periods where the technology was not yet functioning correctly and not yet sending data to the server, rather than true zero consumption hours.

The resulting sample consists of 132,375,282 hourly observations across 5,975 households over 1,186 days. Table 1 provides descriptive statistics for the entire sample as well as for the balanced panel of 4,651 households spanning 2011 and 2012.


Table 1: Descriptive Statistics

	Full Sample	Balanced Panel
Number of Observations	132,375,282	81,597,144
Number of Households	5,975	4,651
Mean Observations Per Household	22,154	17,544
Min kWh Reading	0.00	0.00
Mean Household Min	.018	.027
Mean Household Mean	2.02	2.07
Mean Household Max	12.17	11.71
Max kWh Reading	127.4	127.4
Percentage kWh=0	1.7%	1.1%

Descriptive statistics for full sample from 1 January 2010 - 31 March 2013
and for the balanced panel from 1 January 2011 - 31 December 2012.

Figure 1 presents a sample of two weeks of data for four randomly chosen households. Three types of variation in the data are apparent. The first is intra-day variation: the within-day distribution of electricity consumption appears to be bi-modal, with consumption peaking during the morning and evening hours. The second is cross-sectional variation between households: certain households consistently consume more electricity on average. Lastly, seasonal climate cycles contribute to intra-annual variation. Household electricity use is on average higher in the winter months, consistent with widespread use of electric heating appliances. In addition, it appears that some households increase their electricity consumption in winter more than other households. Departing somewhat from other settings, levels and patterns of electricity consumption are similar on weekends and weekdays.

Figure 1: Two weeks of hourly electricity consumption by four random households



Each panel displays hourly consumption over one week for four randomly selected households. Panel A displays a sample week in summer 2012. Panel B displays a sample week in winter 2010.

4.2 Weather data

I use hourly weather data from five weather stations in Northern Johannesburg, collected by the South African Weather Service (SAWS).³ The maximum distance between any household and any weather station is less than 5 miles, with very little geographic heterogeneity. Since the support across the temperature distribution is equal for all households, selection bias that might occur if certain households choose to locate in areas with certain temperatures is not a concern. To minimize measurement error, all households are assigned the same hourly temperature time series consisting of the average temperature across all five stations. Figure A2 presents the full time series of daily minimum, mean, and maximum temperatures. Temperatures are highest in the South African summer months (December - February) and lowest in winter months (June - September).

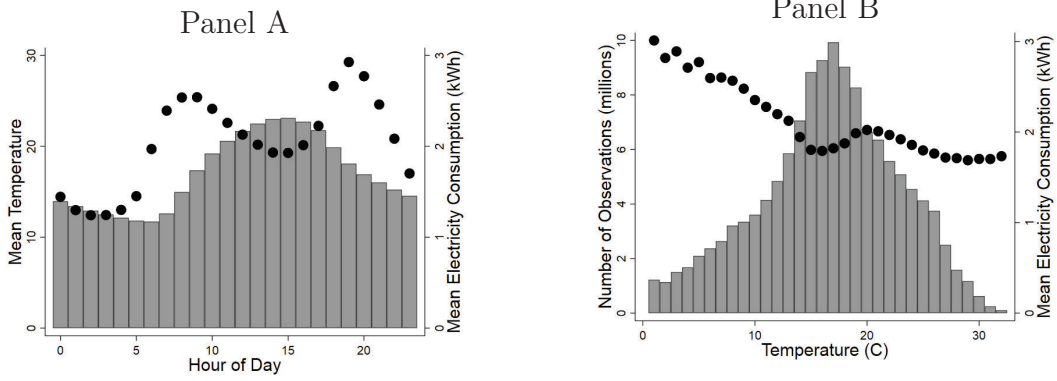
4.3 Descriptive evidence

Daily electricity consumption is largely bi-modal, peaking between 7:00 - 10:00am and between 6:00pm - 8:00pm, when residents are most likely to be at home and awake. This also happens to be when outdoor temperatures are on average above $20^{\circ}C$. Figure 2 displays this correlation. Panel A displays the average hourly temperatures and mean kWh consumption over a single day,

³Temperatures are recorded by the SAWS using advanced, identical pieces of equipment that are certified to operate from $-10^{\circ}C$ to $60^{\circ}C$, well outside the observed temperature range $-3.8^{\circ}C$ to $33.8^{\circ}C$.

and Panel B displays the number of observations in, and the average kWh of electricity consumed per household during, each temperature bin.

Figure 2: Hourly temperature and electricity consumption distributions



Bars represent the mean temperature at each hour. Dots represent average electricity consumption at each hour.

Bars represent the number of observations in each temperature bin. Dots represent average electricity consumption for each temperature bin.

The endogeneity of electricity supply with respect to weather is a possible concern to identification. Weather might affect electricity consumption not only by affecting individual behavior but by affecting electricity supply through technological constraints. The households included in this study face a tariff schedule that is constant over time, with the exception of annual increases largely in line with inflation. This endogeneity therefore cannot affect results through a time-varying pricing channel. There may still be a concern if technological constraints in electricity supply disproportionately cause power outages during specific temperatures, which mechanically would reduce demand during these periods. Figure A1 demonstrates that power outages are very rare in this context.

5 Identification

To estimate the relationship between outdoor temperatures and household electricity consumption, hourly temperatures are assigned to one of 33 temperature bins. Each temperature bin has a width of $1^{\circ}C$, such that each temperature dummy $T_t^j = 1 \forall Temp_t \in [j, j + 1)$ at time t for

temperature bins $j = 1, \dots, 33$. Temperatures below $2^\circ C$ are allocated to the lowest temperature bin ($< 2^\circ C$) and temperatures above $33^\circ C$ to the highest temperature bin ($> 33^\circ C$). Unless stated otherwise, each regressions estimates a separate coefficient for the effect of temperature on household electricity consumption for each temperature bin. This methodology allows for a more flexible estimation of the effects of temperature on electricity consumption, and enables identification of non-linearities in temperature response functions without imposing any modelling restrictions (see for example Engle et al. (1986) for an early discussion on semiparametric estimation of temperature response functions). Regressions exclude temperature bin $28 - 29^\circ C$, hence coefficients for all other temperature bins estimate the marginal effect on electricity consumption of an additional hour in each temperature bin, relative to an hour with temperatures in the omitted category, following Deryugina and Hsiang (2014). To control for the correlation displayed in Figure 2 all three specifications include an *hour of day* fixed effect ϕ_h . To reflect heterogeneous average levels of consumption across households, the dependent variable in all regressions is log of kWh consumed per hour. This excludes all observations with zero kWh consumed, and reduces the sample from 132,375,282 to 130,399,087 hourly observations. Figure A4 confirms that results are consistent when using kWh as the dependent variable and when using the full sample.

The three preferred specifications regress the outcome variable on 33 temperature bins as follows:

$$y_{it} = \sum_{j=1}^{33} (\beta_j T_t^j) + \phi_h + \epsilon_{it} \quad (1)$$

$$y_{it} = \sum_{j=1}^{33} (\beta_j T_t^j) + \phi_{hm} + \epsilon_{it} \quad (2)$$

$$y_{it} = \sum_{j=1}^{33} (\beta_j T_t^j) + \phi_{hw} + \epsilon_{it} \quad (3)$$

where y_{it} is log of kWh consumed by household i at time t . Each coefficient β_j can be interpreted as the causal effect on household electricity consumption of one hour at each temperature bin relative to the $28 - 29^\circ C$ bin, which is excluded from the regression. ϵ_{it} is the idiosyncratic error term.

Specification 2 includes a fixed effect ϕ_{hm} that interacts the *hour of day* fixed effect with a *month of year* fixed effect. Most conservatively, Specification 3 interacts an hour fixed effect by a week-of-sample fixed effect, ϕ_{hw} . These specifications allow for the separate identification of seasonal temperature effects and daily temperature effects. Seasonal adaptation involves households adjusting their appliance usage and settings based on seasonal temperature variation. Daily adaptation to temperature involves households adapting their usage of electric heating or air conditioning to hourly or daily fluctuations in weather. If households use a fixed set of appliances in different seasons, but do not adapt the usage of these appliances to daily weather fluctuations, seasonal adaptation would dominate. Alternatively, if household electricity consumption is flexible and can respond to a range of daily weather fluctuations, regardless of the season, daily adaptation would dominate. Any difference between Specifications 1, 2, and 3 will be driven by whether temperature affects electricity consumption on a daily versus seasonal basis.


The robustness checks in Figure A5 include additional fixed effects, including household fixed effects, weekend fixed effects, and year fixed effects, and their interactions. Including all of these additional fixed effects does not change the above coefficients meaningfully. Most meaningful variation is captured in the three specifications above.

Standard errors are clustered conservatively, by household and by week of sample, to control for exogenous shocks that are correlated across all households such as holidays or public events. Figure A7 presents regression results for more and less conservative levels of clustering.

6 Results

I estimate the coefficients $\beta_j : j = 1, \dots, 32$ from Specifications 1, 2, and 3 via OLS using 130,399,087 hourly observations from 5,975 households during the period 1 January 2010 - 31 March 2013. Estimated coefficients and 95% confidence intervals for each temperature bin are plotted in Figure 3 below.

Figure 3: Regression coefficients for 32 temperature bin dummies



Coefficient estimates for each of 32 temperature bin dummies. Panel A includes hour of day fixed effects only. Panel B includes hour of day by month of year fixed effects. Panel C includes hour of day by week of sample fixed effects. Figure A5 includes additional fixed effects that do not change the results meaningfully. Standard errors are clustered by household and by week of sample. Shaded regions are 95% confidence intervals. Omitted category is 28-29°C.

Across all three specifications, household electricity consumption declines approximately log-linearly across the entire range of 2 – 23°C, remains constant between temperatures in the range 23 – 30.5°C, and then appears to increase sharply at temperatures > 30.5°. Including a fixed effect ϕ_{hm} that interacts the *hour of day* fixed effect with a *month of year* fixed effect in Specification 2 flattens the negative slope among low and mild temperatures, but the negative log-linear relationship remains. Including a fixed effect ϕ_{hw} that interacts the *hour of day* fixed effect with a *week-of-sample* fixed effect in Specification 3 flattens the slope further. This suggests that hourly responses to short-term changes in temperatures are not as pronounced as responses to intra-year seasonal changes in average temperatures. This difference could be driven either by household behavior, where usage of electrical appliances is relatively fixed over shorter periods, or by a difference between outdoor and indoor temperature caused by the trapping of heat inside buildings, causing insulation to dampen the short-term effect of outdoor temperatures on indoor electricity consumption. Unfortunately, these data do not allow for differentiation between these mechanisms.

Table A1 presents numerical regression results. Figure A3 presents results using a balanced panel from January 1, 2011 - December 31, 2012. Including only the balanced panel does not meaningfully affect the results.

6.1 Log-linear relationship

Despite estimating a coefficient for each temperature bin individually, which does not assume linearity, the results suggest a strong and robust negative log-linear relationship between temperature and electricity consumption below a heating threshold and a sharp positive relationship above a cooling threshold. Following Burke and Emerick (2016) and Waite et al. (2017) this analysis estimates the following piecewise linear function, where T_L represents the lower threshold, for heating, and T_H represents the higher threshold, for cooling:

$$y_{it} = \beta_1 Temp_t + \beta_2 Temp_{t;T < T_L} + \beta_3 Temp_{t;T > T_H} + \phi_h \quad (4)$$

where y_{it} is log of kWh consumed by household i at time t and $Temp_t$ is temperature at time t . The variable $Temp_{t;T < T_L} = (T_L - Temp_t) \cdot \mathbb{1}(Temp_t < T_L)$ is the difference between temperature and the heating threshold T_L interacted with an indicator variable for temperature being lower than T_L . $Temp_{t;T > T_H}$ is similarly defined for temperatures higher than the cooling threshold. In Specification 4, β_1 can be interpreted as the percentage increase in kWh for every $1^\circ C$ increase in temperature inside the interval (T_L, T_H) , $(\beta_2 - \beta_1)$ can be interpreted as the percentage increase in kWh for every $1^\circ C$ decrease in temperature below T_L , and $(\beta_3 + \beta_1)$ can be interpreted as the percentage increase in kWh for every $1^\circ C$ increase in temperature above T_H . As before, ϕ_h is an hour of day fixed effect.

An algorithm searches across heating thresholds $21 - 29^\circ C$ and cooling thresholds $25 - 35^\circ C$ in 0.5 degree intervals to identify the change-point that minimizes the Root Mean Squared Error (RMSE). The RMSE is minimized for a heating threshold of $23^\circ C$ and a cooling threshold of $30.5^\circ C$. This is consistent with increased usage of air conditioning above $30^\circ C$, when fans and other less energy-intensive cooling technologies are no longer sufficient, and is approximately consistent with previous literature. For example, Waite et al. (2017) perform a meta-analysis of macroeconomic and engineering analyses of household temperature response functions across 36 cities in Africa, Asia, and the U.S., and find that thresholds generally fall within the $15 - 25^\circ C$ range, with non-

OECD cities generally having higher thresholds than OECD countries. The heating threshold of $23^{\circ}C$ estimated above indeed falls in the upper end of their range.

The results to Specification 4 using the temperature thresholds of $23^{\circ}C$ and $30.5^{\circ}C$ are presented in Table 2 below:

Table 2: Coefficient estimates across temperature bins

	(1)	(2)	(3)
$Temp_t$	-0.0134*** (0.0023)	-0.0054*** (0.0015)	-0.0088*** (0.0011)
$Temp_{t;T < T_L}$	0.0279*** (0.0026)	0.0123*** (0.0023)	0.0014 (0.0018)
$Temp_{t;T > T_H}$	0.0943*** (0.0187)	0.0541*** (0.0154)	0.0406** (0.0130)
Observations	130,399,087	130,399,087	130,399,087
Hour Fixed Effects	Yes	Yes	Yes
Hour X Month FE	No	Yes	No
Hour X Week of Sample FE	No	No	Yes

Dependent variable is log of kWh. Temperature in $^{\circ}C$ interacted with dummies for Low and High temperature according to Specification 4. Column (1) includes hour of day fixed effects only. Column (2) includes hour of day by month of year fixed effects. Column (3) includes hour of day by week of sample fixed effects. Standard errors clustered by household and by week of sample. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

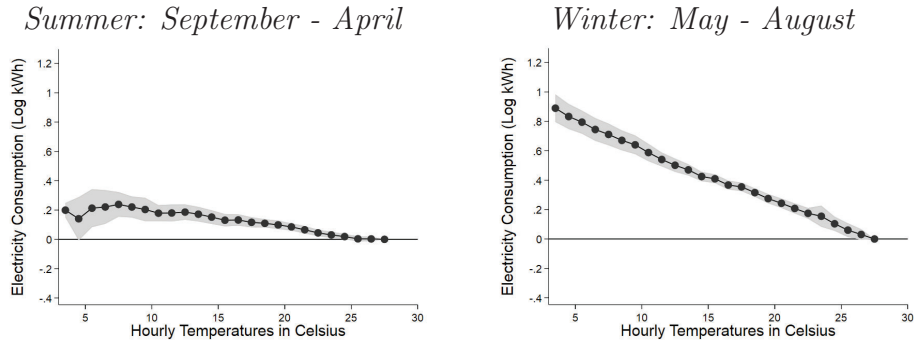
The coefficients in Column (1), which do not control for seasonal variation, indicate a 4.1% increase in electricity consumption for every $1^{\circ}C$ below the heating threshold $T_L = 23$, and a 12.2% increase in electricity consumption for every $1^{\circ}C$ above the cooling threshold $T_H = 30.5$. The coefficients in Column (2), which controls for month of year, represent a much smaller 1.7% increase in electricity consumption for every $1^{\circ}C$ below the heating threshold and a much smaller 6.6% increase in electricity consumption for every $1^{\circ}C$ above the cooling threshold. Finally, the coefficients in Column (3), which controls for week of sample, represent a 1.0% increase in electricity consumption for every $1^{\circ}C$ below the heating threshold and a 4.2% increase in electricity consumption for every $1^{\circ}C$ above the cooling threshold. Combined, this evidence suggests that electrical appliance usage responses to both heating and cooling temperatures are more strongly

driven by seasonal variation than hourly variation in temperatures.

6.2 Seasonal variation

Figure 3 demonstrates that the within-week electricity consumption response to hourly temperature fluctuations differs substantially from the response to seasonal temperatures. To provide more detail on how short-term temperature responsiveness changes seasonally, Figure 4 presents Specification 1 using data from summer months and winter months separately.

Figure 4: Hourly kWh consumption temperature bin coefficients by season



Coefficient estimates for each of 32 temperature bin dummies separated by summer and winter months. Specification includes hour of day fixed effects. Standard errors clustered by household. Shaded regions are 95% confidence intervals. Omitted category is 28-29°C.

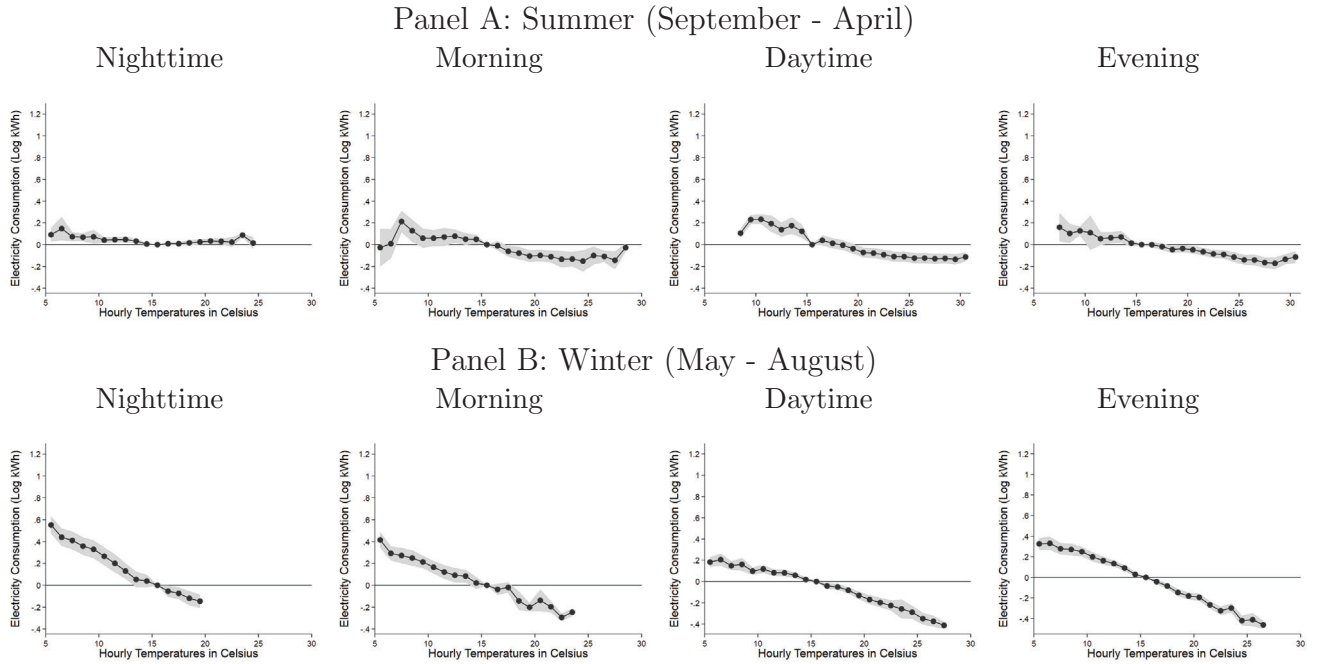
Households display almost no temperature response throughout the spring, summer, and fall months of September through April. The electricity response function with respect to temperature in each of those months appears much less steep. In contrast, household electricity consumption displays a steep negative response function during the winter months of May through August, which is likely driving the main results seen in Figure 3. This suggests households respond to cold temperatures in winter on a daily basis by employing an appliance such as an electric heater, but do not take similar short-term response measures to low temperatures during relatively warmer months.

6.3 Variation across hour of the day

Temperature changes might not cause significant changes in energy consumption between 10pm - 6am, when most household members are asleep. On the other hand, household preferences over

different temperatures may vary over hours of the day, causing differential appliance use responses. To investigate whether temperature responses differ significantly by hour of day, Figure 5 presents regression results for Specification 3 using data from winter and summer and from four separate times of day separately: *nighttime* (10pm - 6am), *morning* (6am - 10am), *daytime* (10am - 5pm), and *evening* (5pm - 10pm). The positive effect of low temperatures on electricity consumption in winter is consistent across all times of day, while the positive effect of high temperatures on electricity consumption is driven by nighttime temperatures in summer. This indicates that during winter, residents are sensitive to low temperatures across all times of day.

Figure 5: Hourly kWh usage temperature bin coefficients by time of day and season



Coefficient estimates for each of 32 temperature bin dummies. Specification includes hour of day fixed effects. Standard errors clustered by household and by week of sample. Shaded regions are 95% confidence intervals. Due to sample size restrictions the omitted category is 15-16°C and temperature bins are restricted to between 5-31°C. Times of day are defined as *nighttime* (10pm - 6am), *morning* (6am - 10am), *daytime* (10am - 5pm), and *evening* (5pm - 10pm).

6.4 Daily responses


A number of papers (Auffhammer (2018), Davis and Gertler (2015), and Gupta (2014), among others), estimate temperature coefficients using daily data. To compare these results to their findings, Specification 5 estimates the relationship between temperature and electricity consumption

in a manner analogous to the econometric specifications in the previous section but using total daily electricity consumption and average daily temperature instead.

$$y_{id} = \sum_{j=1}^{33} (\beta_j T_d^j) + \epsilon_{it} \quad (5)$$

where y_{id} is log of kWh consumed by household i on day d . T_d^j ($j = 1, \dots, 33$) are dummies for 33 unique temperature bins of width $1^\circ C$, each representing average temperature on day d . Figure 6 presents results from Specification 5 using no fixed effects and month of year fixed effects.

Figure 6: Effect of daily temperature on daily electricity consumption



Coefficient estimates for each of 24 temperature bin dummies. Due to small samples at higher temperatures for daily averages, omitted category is any day with mean temperature $> 24^\circ C$. Panel A does not include fixed effects. Panel B includes month of year fixed effects. Panel C includes week of sample fixed effects. Figure A6 includes additional fixed effects that do not change the results meaningfully. Standard errors clustered by household and by month of sample. Shaded regions are 95% confidence intervals.

Results using daily data are consistent with the previous analysis using hourly data. Both specifications estimate a negative relationship between temperature and electricity use. Panel A displays a strong negative log-linear slope between $2 - 23^\circ C$. The results in Panel B and Panel C indicate a more muted but consistent negative log-linear relationship over this same interval. This suggests that in winter, when days consistently average below $15^\circ C$, households respond by consistently using their heating appliances. When controlling for seasonal variation, households are much less responsive to daily temperature fluctuations. There are not enough days with average temperatures higher than $25^\circ C$ to identify a positive slope above a threshold temperature, as identified using hourly data in the previous section.

These results differ significantly from a number of previous findings in the literature. Auffhammer (2018) estimates 1,235 sets of temperature dummy coefficients for each of 1,235 zip

codes across California. At the $25 - 26^{\circ}C$ bin, the 10-90% confidence interval of household electricity consumption responsiveness to temperature changes ranges from $0.3 - 2\%$ relative to the $15 - 16^{\circ}C$ bin. The estimates presented here instead demonstrate a consistent *negative* average effect over this interval. This is reasonable when considering the much lower rate of air conditioning penetration in South Africa relative to California. Davis and Gertler (2015) conduct a similar exercise using monthly billing data from Mexico and find that households on average consume constant electricity between $10 - 24^{\circ}C$, consistent with lower usage of electric heating and higher usage of gas for heating in Mexico than in South Africa.

Specifications using daily minimum or daily maximum temperature respectively instead of daily mean temperature display the same consistent negative relationship between household electricity use and temperature throughout the temperature distribution. Households are especially responsive when the minimum daily temperature drops below $7^{\circ}C$.

6.5 Robustness checks

Figure A5 demonstrates that results are consistent along an extensive set of fixed effect specifications. Regardless of any additional fixed effects included in the specification, any regression that includes hour of day but not month or week of sample fixed effects closely resembles results from Specification 1; any regression that includes hour of day fixed effects interacted with month of year fixed effects closely resembles results from Specification 2, and any regression that includes hour of day fixed effects interacted with week of sample fixed effects closely resembles results from Specification 3. This suggests that these dimensions drive the results. Robustness checks for sensitivity to fixed effect specifications for regressions using daily data, presented in Figure A6, confirm that those results are also robust.

If power outages or missing meter measurements were correlated with temperatures, which is technologically feasible, this may bias the results. To test for this, Figure A4 presents results to estimation of the main specification using kWh consumption instead of log of kWh consumption

as the dependent variable and an estimation that excludes observations where hourly consumption was 0 kWh (as reported in Table 1, this constitutes 11.5% of observations). Results are consistent with the main results.

Figure A7 confirms that standard errors do not change significantly and the main results retain statistical significance at the 95% confidence level across varying levels of clustering.

Finally, extensive robustness checks that include hourly and daily temperature lags as well as dummies for cold or hot spells of consecutive hours and days of low or high temperatures respectively do not reveal systematic patterns.

7 Simulation


Climate change is expected to cause significant increases in temperatures in the coming decades. This analysis uses the causal relationship identified between temperature and electricity consumption to project electricity consumption under a new, increased distribution of temperatures. While other attributes, such as appliance ownership and average energy efficiency, are likely to change significantly and thereby meaningfully affect energy consumption over the coming decades, for the sake of clarity this simulation holds constant all determinants of electricity consumption other than temperature.

The effect of climate change on weather in South Africa has been studied widely. Future rainfall predictions rely on uncertain projections for ocean temperatures, humidity, the strength of the jet stream, and other physical characteristics in the Southern Africa region. There is therefore a significant amount of uncertainty regarding future rainfall patterns. On the other hand, the Climate Change Adaptation plan by the Department of Agriculture, Forestry and Fisheries, Republic of South Africa (2015) makes a precise prediction on future temperature changes. By mid-century, the coastal regions of South Africa are expected to experience an average increase in temperatures of $1.5 - 2.5^{\circ}\text{C}$, while the interior regions are expected to experience an average increase in temperatures of $3.0 - 3.5^{\circ}\text{C}$. Given Johannesburg's location in the interior of South

Africa, I use an expected increase of 3.25°C to project future consumption. Appendix Figure A8 presents equivalent results for temperature increases of 1.25°C and 5.25°C .

Figure 7 displays average electricity consumption in each temperature bin as well as the current and projected temperature distributions. The projected temperature distribution is defined as current temperatures shifted upwards by 3.25°C . The highest bin includes all observations greater than 33°C .

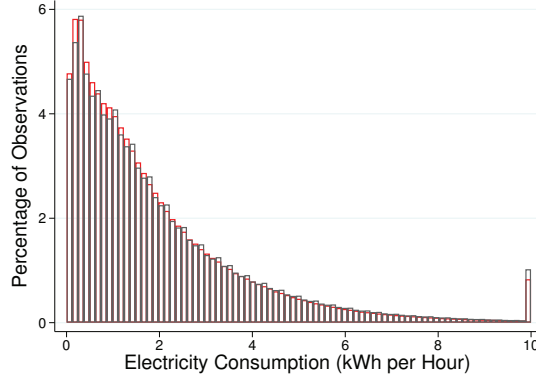
Figure 7: Electricity consumption per bin overlaid with the current and projected temperature distributions



Graph displays electricity consumption and temperature distribution simultaneously. Bars represent the fraction of observations in each temperature bin (values on left y -axis). Grey bars represent the current temperature distribution while red bars represent the projected future temperature distribution. Dots represent average consumption in kWh for each temperature bin (values on right y -axis).

I store the regression estimates from my main specification, replace the temperature values with the shifted temperature values, and apply the estimated coefficients to the new temperature values to project the new electricity consumption estimates. By including the residuals from the original regression to predict consumption, this analysis identifies the marginal effect of the projected temperature change. It is important to acknowledge once again that the projected effects are only on the *intensive* margin. Any reading where 0 kWh was consumed will be predicted to have 0 kWh of consumption even at higher temperatures. Figure 8 below presents the observed and the predicted distribution of kWh readings under the observed and projected temperature distributions.

Figure 8: Histogram of hourly kWh readings by sample overlaid with projected consumption



Graph displays percentage of electricity consumption reading in each kWh bin of width 0.25. Grey bars represent the true data. Red bars represent the predicted values using the shifted temperature distribution.

Figure 8 suggests that the increased temperature distribution would shift left relative to current consumption. The average annual electricity consumption per household decreases from 17.6 MWh (median 15.7 MWh) under the current temperature distribution to 16.6 MWh (median 14.7 MWh) under the projected temperature distribution. These results suggest that an increase in temperatures of $3.25^{\circ}C$ would cause an overall decrease of 1,093.4 kWh per year per household, equivalent to 6.2% relative to baseline consumption.

8 Conclusion

This paper uses 132,375,282 hourly electricity consumption data from 5,975 households located in Johannesburg, South Africa to identify the causal effect of short-term temperature changes on household electricity consumption. These data provide a unique opportunity to study household electricity consumption patterns in an area of the world where this has not been studied in such detail previously, but where energy usage is expected to increase significantly over the next several decades.

A flexible estimation of the effect of temperature on electricity consumption using temperature bins of width $1^{\circ}C$ reveals a consistent negative log-linear relationship. An algorithm searching

across the temperature distribution in $0.5^{\circ}C$ intervals to minimize the model’s Root Mean Squared Error (RMSE) identifies a heating threshold of $23^{\circ}C$ and an cooling threshold of $30.5^{\circ}C$. Across the range of $2 - 23^{\circ}C$, every $1^{\circ}C$ increase in temperature causes a reduction in household electricity consumption of approximately 4.1%. While there is a limited number of observations at temperatures above $30^{\circ}C$, there is suggestive evidence that electricity consumption increases steeply when temperatures exceed the air conditioning use threshold of $30.5^{\circ}C$. Detailed analyses of these results over time of day and month of year reveal that this result is primarily driven by seasonal usage of electric heating during winter. A large portion of household responses to temperature changes occur over seasonal temperature changes. In the short-term, households are more responsive to low temperature in winter than in summer, suggesting that they employ electric heating on cold days in the winter, but do not increase electricity consumption similarly by employing electric heating during cold days in the summer. These results are robust to a large number of fixed effect specifications, different levels of clustering, and sub-sample analyses to account for outliers.

Using these results, the analysis then simulates household electricity consumption under an upward shift of the temperature distribution of $3.25^{\circ}C$. Under this scenario, aggregate household electricity consumption would decrease by 1,093.4 kWh per year per household, equivalent to 6.2% relative to baseline consumption. This is a departure from previous literature (Deschenes and Greenstone (2011), Davis and Gertler (2015), and Dyson et al. (2014), among others) that identifies a positive effect of temperature on household energy consumption at similar temperatures. Policy makers would benefit from future work that identifies how specific underlying differences, such as income, geography, infrastructure, or tariff structures, drive these diverging results. Regardless of the specific underlying drivers, the results presented in this paper suggest that there is significant heterogeneity in the way temperatures affect household energy demand across the world.

This result has meaningful implications for policymakers. Mornings and evenings experience both lower temperatures than hours in the middle of the day, as well as increased demand resulting from most people’s daily movement patterns. Energy planners working to limit out-

ages due to capacity constraints will need to incorporate both drivers into their forecasts. Policy makers looking to reduce the need for capacity increases by reducing demand through energy efficiency improvements would benefit from targeting their energy efficiency programs towards electric heating appliances. As growing middle-class incomes in developing countries cause increased purchases of both heating and cooling appliances, it will be important to encourage energy efficiency improvements not only among air conditioners but among electric heaters, as well.

This research does not study the extent to which power outages and other supply-side mechanisms influence the relationship between temperature and household electricity consumption, focusing instead on demand-side responses. In addition, the analysis is limited to responses to temperature and does not factor in additional weather variation such as precipitation, cloudiness, or humidity. These questions provide opportunity for follow-up work.

References

- The South African National Census of 2011. *Statistics South Africa*. Available at: <https://www.census2011.co.za/>.
- Aroonruengsawat, A. and Auffhammer, M. (2011). Impacts of climate change on residential electricity consumption: Evidence from billing data. *The Economics of Climate Change: Adaptations Past and Present*, pages 311–342.
- Auffhammer, M. (2014). Cooling China: The weather dependence of air conditioner adoption. *Frontiers of Economics in China*, 9(1):70–84.
- Auffhammer, M. (2018). Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption using big data. NBER Working Paper #24397.
- Auffhammer, M., Baylis, P., and Hausman, C. (2016). Climate change is projected to have severe

- impacts on the frequency and intensity of peak electricity demand across the United States. *Proceedings of the National Academy of Sciences*, 114(8):18861891.
- Burke, M. and Emerick, K. (2016). Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy*, 3(8):106–140.
- Chikobvu, D. and Sigauke, C. (2010). Daily peak electricity load forecasting in South Africa using a multivariate non-parametric regression approach. *The Operations Research Society of South Africa*, 26(2):97–111.
- Chikobvu, D. and Sigauke, C. (2013). Modelling influence of temperature on daily peak electricity demand in South Africa. *Journal of Energy in Southern Africa*, 24(4).
- Coelho, S. T., Lecocq, F., Cortez, C., and Tudeschini, L. G. (2014). Fuel wood consumption in Brazilian residential sector, energy consumption in households and carbon footprint of development in selected Brazilian regions, comparing Brazil and France. *European Biomass Conference and Exhibition*, 22.
- Crowley, C. and Joutz, F. L. (2003). Hourly electricity loads: Temperature elasticities and climate change. *23rd US Association of Energy Economics North American Conference*.
- Davis, L. W. and Gertler, P. J. (2015). Contribution of air conditioning adoption to future energy use under global warming. *Proceedings of the National Academy of Sciences of the United States of America*, 112:5962–5967.
- Department of Agriculture, Forestry and Fisheries, Republic of South Africa (2015). Draft climate change adaptation and mitigation plan for the South African agricultural and forestry sectors.
- Deryugina, T. and Hsiang, S. M. (2014). Does the environment still matter? Daily temperature and income in the United States. NBER Working Paper #20750.

- Deschenes, O. and Greenstone, M. (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics*, 3:152–185.
- Dunbabin, P., Palmer, J., and Terry, N. (2015). Electricity use by domestic appliances in English households. *European Council for an Energy Efficient Economy*.
- Dyson, M., Borgeson, S. D., Michaelangelo, D. T., and Callaway, D. S. (2014). Using smart meter data to estimate demand response potential, with application to solar energy integration. *Energy Policy*, 73:607–619.
- Ekouevi, K. and Tuntivate, V. (2011). Household energy access for cooking and heating: Lessons learned and the way forward. *World Bank Energy and Mining Sector Board Discussion Paper*.
- Engle, R. F., Granger, C. W. J., Rice, J., and Weiss, A. (1986). Semiparametric estimates of the relation between weather and electricity sales. *Journal of the American Statistical Association*, 81(394):310–320.
- Garland, R. M., Matooane, M., Engelbrecht, F. A., Bopape, M.-J. M., Landman, W. A., Naidoo, M., van der Merwe, J., and Wright, C. Y. (2015). Regional projections of extreme apparent temperature days in Africa and the related potential risk to human health. *International Journal of Environmental Research and Public Health*, 12:12577–12604.
- Gupta, E. (2014). The effect of development on the climate sensitivity of electricity demand in India. *Discussion Papers in Economics*, 14(05).
- Jack, B. K. and Smith, G. (2015). Pay as you go: Prepaid metering and electricity expenditures in South Africa. *American Economic Review: Papers and Proceedings*, 5(105):237–241.
- Jack, B. K. and Smith, G. (2016). Charging ahead: Prepaid electricity metering in south africa. NBER Working Paper #22895.

- Petri, Y. and Caldeira, K. (2015). Impacts of global warming on residential heating and cooling degree-days in the United States. *Scientific Reports*, 5(12427).
- Shah, N., Wei, M., Letschert, V., and Phadke, A. (2015). Benefits of leapfrogging to superefficiency and low global warming potential refrigerants in room air conditioning. *Ernest Orlando Lawrence Berkeley National Laboratory*.
- U.S. Energy Information Administration (EIA) (2015). Residential energy consumption survey.
- Waite, M., Cohen, E., Torbey, H., Piccirilli, M., Tian, Y., and Modi, V. (2017). Global trends in urban electricity demands for cooling and heating. *Energy*, 127:786–802.
- Wang, H. and Chen, Q. (2014). Impact of climate change heating and cooling energy use in buildings in the United States. *Energy and Buildings*, 82:428–436.
- Wenz, L., Levermann, A., and Auffhammer, M. (2017). North-south polarization of European electricity consumption under future warming. *Proceedings of the National Academy of Sciences*, pages E7910–E7918.
- Wolfram, C., Shelef, O., and Gertler, P. (2012). How will energy demand develop in the developing world? *Journal of Economic Perspectives*, 26(1):119–138.


Appendix

Figures

Figure A1 plots the proportion of observations in a given day during which a particular household consumes 0 kWh. The figure suggests that the high-frequency electricity consumption monitoring equipment was turned on in a staggered manner across households. For this reason, all the analyses in this paper exclude all 0 kWh observations that occur prior to the first positive consumption observation. This mechanically also results in dropping all households for which all consumption values are 0 kWh.


If outages were a common occurrence, especially in a geographically small area such as this sample, we would expect to see specific days during which entire groups of households consumed 0 kWh during multiple hours. This would be represented by vertical dark lines in Figure A1. This pattern largely does not appear in the data. This confirms anecdotal evidence by South African economists and utility employees that power outages are rare in this area of Johannesburg.

Figure A1: Proportion of 0 kWh consumption observations




Each pixel represents the fraction of hours during which each household consumed 0 kWh on that day. Vertical group of darker shading would indicate either a power outage affecting a large number of customers on that day or systematically missing data. Such incidences do not appear to be common. Households are ordered by the date on which they entered the sample.

Figure A2: Temperature time-series



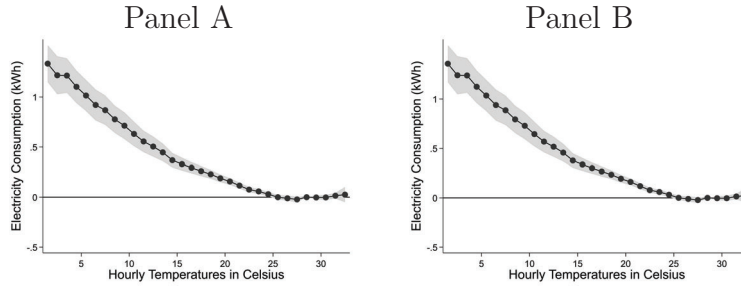
Daily minimum, mean, and maximum temperature averaged across the five weather stations for the full sample period, 1 January 2010 through 31 March 2013.

Figure A3: Regression coefficients using balanced panel



Coefficient estimates for each of 32 temperature bin dummies restricting the sample to a balanced panel of 4,651 households from 1 January, 2011 - 31 December, 2012. Panel A includes hour of day fixed effects only. Panel B includes hour of day by month of year fixed effects. Panel C includes hour of day by week of sample fixed effects. Standard errors clustered by household and week of sample. Shaded regions are 95% confidence intervals. Omitted category is 28-29°C. Including only a balanced panel does not affect results relative to those presented in Figure 3.

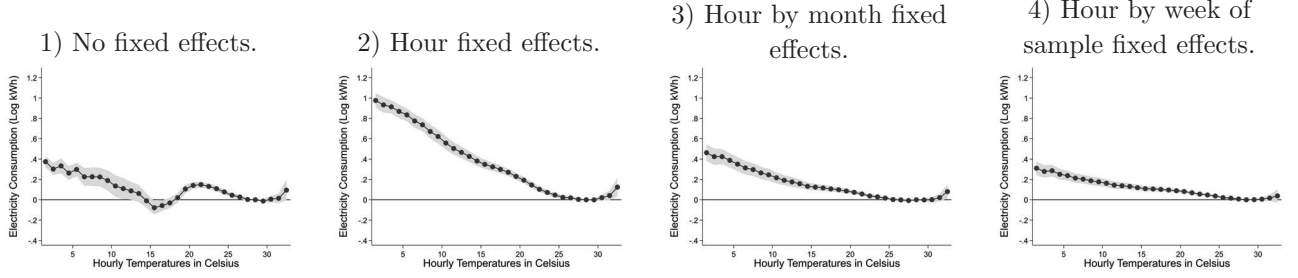
Figure A4: Robustness checks: Log scale



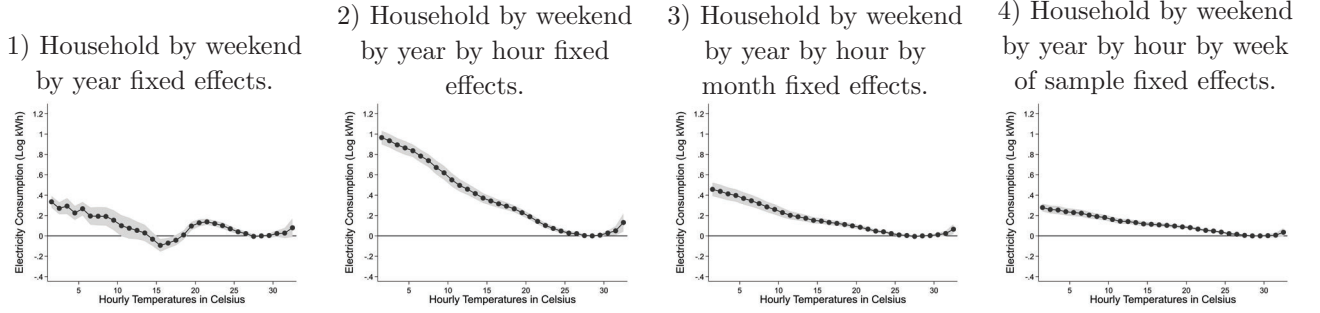
Coefficient estimates for each of 32 temperature bin dummies using kWh as the outcome variable, rather than log(kWh). Panel A includes all observations. Panel B excludes observations where the outcome variable equals 0. Both specifications include hour by month fixed effects. Standard errors clustered by household and by week of sample.

Figure A5: Robustness checks: Fixed effects

Panel A: Basic fixed effects only




Panel B: Basic fixed effects interacted with household by weekend by year fixed effects



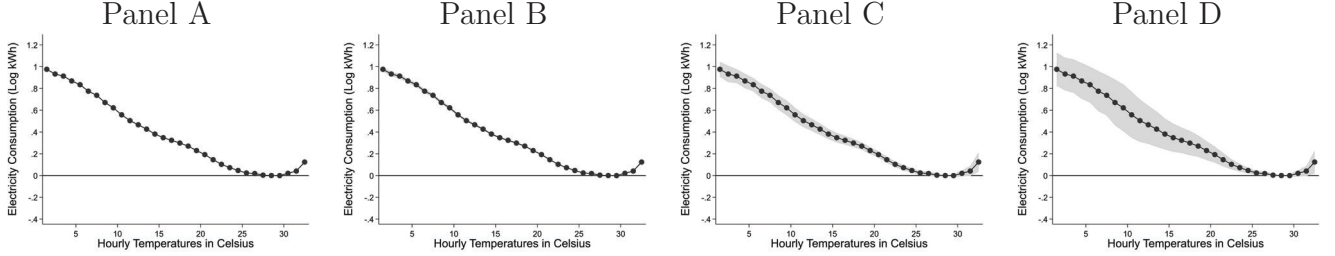
Coefficient estimates for each of 32 temperature bin dummies. Standard errors clustered by household and by week of sample. Shaded regions are 95% confidence intervals. Omitted category is 28-29°C. Interacting fixed effects with household, weekend, and year fixed effects does not affect results relative to those presented in Figure 3.

Figure A6: Robustness checks: Fixed effects for daily responsiveness



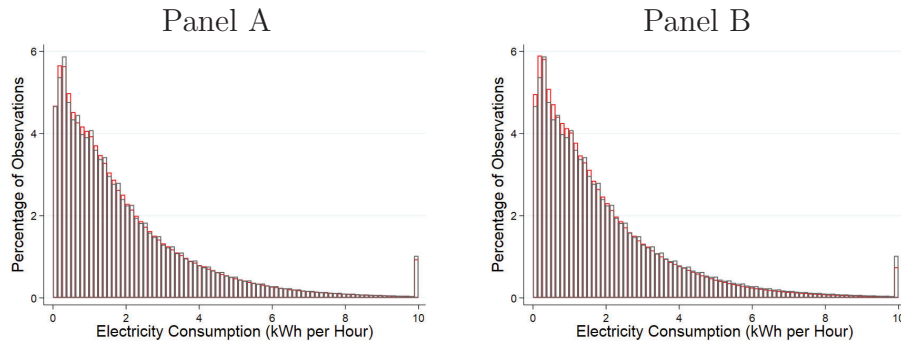
Coefficient estimates for each of 24 temperature bin dummies. Due to small samples at higher temperatures for daily averages, omitted category is any day with mean temperature $> 24^{\circ}\text{C}$. Panel A includes year by household by weekend fixed effects. Panel B includes year by household by weekend by month of year fixed effects. Panel C includes year by household by weekend by week of sample fixed effects. Standard errors clustered by household and month of sample. Shaded regions are 95% confidence intervals. Interacting fixed effects with household, weekend, and year fixed effects does not affect results relative to those presented in Figure 6.

Figure A7: Robustness checks: Clustering



Graphs present coefficient estimates from Specification 1 for each of 33 temperature bin dummies, as well as bars representing 95% confidence interval. All specifications include hour fixed effects. Standard errors in the regression in Panel A are not clustered. Standard errors in the regression in Panel B are clustered by household only. Standard errors in the regression in Panel C are two-way clustered by household and week of sample. Standard errors in the regression in Panel D are three-way clustered by household, week of sample, and month of year.

Figure A8: Histogram of hourly kWh readings by sample overlaid with projected consumption



Graph displays percentage of electricity consumption reading in each kWh bin of width 0.25. Black bars represent the true data from the random sample of 500 households from years 2011 and 2012. Red bars represent the predicted values using the shifted temperature distribution. Panel A projects usage under a temperature increase of 1.25°Celsius and Panel B projects usage under a temperature increase of 5.25°C.

Tables

Table A1: Coefficient estimates across temperature bins

	(1)		(2)		(3)	
< 2C	0.98***	(0.04)	0.46***	(0.04)	0.31***	(0.03)
2-3C	0.93***	(0.04)	0.42***	(0.04)	0.28***	(0.03)
3-4C	0.91***	(0.03)	0.42***	(0.04)	0.29***	(0.03)
4-5C	0.87***	(0.03)	0.39***	(0.04)	0.25***	(0.03)
5-6C	0.83***	(0.03)	0.35***	(0.03)	0.24***	(0.02)
6-7C	0.77***	(0.03)	0.31***	(0.03)	0.21***	(0.02)
7-8C	0.74***	(0.03)	0.30***	(0.03)	0.20***	(0.02)
8-9C	0.67***	(0.03)	0.26***	(0.03)	0.19***	(0.02)
9-10C	0.62***	(0.03)	0.25***	(0.03)	0.18***	(0.02)
10-11C	0.56***	(0.03)	0.22***	(0.03)	0.16***	(0.02)
11-12C	0.50***	(0.03)	0.19***	(0.02)	0.14***	(0.02)
12-13C	0.47***	(0.03)	0.18***	(0.02)	0.14***	(0.02)
13-14C	0.43***	(0.03)	0.16***	(0.02)	0.13***	(0.02)
14-15C	0.38***	(0.02)	0.13***	(0.02)	0.12***	(0.01)
15-16C	0.35***	(0.02)	0.12***	(0.02)	0.11***	(0.01)
16-17C	0.32***	(0.02)	0.12***	(0.01)	0.11***	(0.01)
17-18C	0.30***	(0.02)	0.11***	(0.01)	0.10***	(0.01)
18-19C	0.27***	(0.02)	0.10***	(0.01)	0.10***	(0.01)
19-20C	0.23***	(0.02)	0.09***	(0.01)	0.09***	(0.01)
20-21C	0.19***	(0.01)	0.07***	(0.01)	0.08***	(0.01)
21-22C	0.15***	(0.01)	0.06***	(0.01)	0.07***	(0.01)
22-23C	0.10***	(0.01)	0.04***	(0.01)	0.06***	(0.01)
23-24C	0.07***	(0.01)	0.03**	(0.01)	0.05***	(0.01)
24-25C	0.05***	(0.01)	0.02*	(0.01)	0.04***	(0.01)
25-26C	0.02*	(0.01)	0.00	(0.01)	0.02***	(0.01)
26-27C	0.02*	(0.01)	-0.00	(0.01)	0.02**	(0.00)
27-28C	0.00	(0.01)	-0.01	(0.01)	0.01	(0.00)
29-30C	-0.00	(0.01)	-0.00	(0.01)	0.00	(0.00)
30-31C	0.02	(0.01)	0.00	(0.01)	0.01	(0.01)
31-32C	0.04	(0.02)	0.02	(0.02)	0.02	(0.01)
> 32C	0.12**	(0.05)	0.08*	(0.04)	0.04	(0.03)
Observations	130,399,087		130,399,087		130,399,087	
Hour FE	Yes		Yes		Yes	
Hour X Month FE	No		Yes		No	
Hour X WOS FE	No		No		Yes	

Coefficient estimates for each of 32 temperature bin dummies from the three main specifications. Outcome variable is log of kWh consumed per hour. Column (1) includes hour of day fixed effects only. Column (2) includes hour of day by month of year fixed effects. Column (3) includes hour of day by week of sample fixed effects. Standard errors clustered by household and week of sample. Omitted category is 28-29°C. Results are presented graphically in Figure 3. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.