Stretching the Duck's Neck: The effect of climate change on future electricity demand

Rivers, Nicholas and Shaffer, Blake

University of Ottawa, University of Calgary

2018

Stretching the duck's neck: The effect of climate change on future electricity demand

Nicholas Rivers[*] Blake Shaffer[†]
University of Ottawa University of Calgary

June 12, 2018

Abstract

This paper examines how climate change will affect both the level and timing of future electricity demand across Canada. Using an original dataset of hourly electricity demand across all Canadian provinces combined with household-level microdata on air conditioner ownership, we estimate temperature responsiveness including both the *direct effect* of temperature on demand for cooling services, as well as the *indirect effect* of increasing the stock of temperature-sensitive durables, such as air conditioners. We find only a small increase in total demand by end-century, although the result differs across provinces. The small aggregate result reflects the mitigating effect of rising temperature in a cold country such as Canada, whereby increases in electricity demand for air conditioning as summer temperatures rise is largely offset by reduced winter heating demand. Although we project limited change in overall electricity demand, we do project changes in the timing of demand, both seasonally and diurnally. In particular, we find seasonal peaks shift from winter to summer in most regions, as well as a large increase in intraday ramping requirements—the difference between minimum and maximum demand within a day—suggesting electricity systems of the future will place an even greater value on storage and flexibility.

[*]Rivers: Graduate School of Public and International Affairs, Institute of the Environment, University of Ottawa (nicholas.rivers@uottawa.ca)

[†]Shaffer: Department of Economics, University of Calgary (blake.shaffer@ucalgary.ca)

1 Introduction

Climate change will affect many social and economic outcomes. The magnitude, and in some cases the sign of these effects, however, remains at the root of considerable debate. Thus, the need for more evidence-based empirical estimates of the potential effects of climate change is important to improve adaptation capacity as well as for understanding its costs.[1] This paper examines the effect of climate change on one such outcome: future electricity demand. We consider how rising temperatures by mid- and end-century will alter both the total amount and timing of electricity demand across Canada. The latter is of particular importance in electricity where supply must equal demand in every hour and storage is costly.

We find a relatively small increase in the aggregate level of demand across Canada, roughly 3.6% even when including the effect of adaptation in the form of more air conditioner penetration and at the most extreme temperature scenario. This stands in contrast to much larger projected increases in future electricity demand from studies in warmer climates (Isaac and van Vuuren, 2009; Akpinar-Ferrand and Singh, 2010; Davis and Gertler, 2015). This result highlights a mitigating effect of a warming climate in a cold country such as Canada, whereby increases in summer cooling demand are largely offset by decreased electric heating demand in the winter.

We find significant heterogeneity in temperature responsiveness across provinces and, correspondingly, regional differences in projected demand changes. This is most notable in projected changes to peak demand—the maximum hourly demand within a year. The largest increase occurs in Ontario, a province that is currently summer-peaking, where we project peak demand to increase by 35%. Whereas, in winter-peaking provinces, such as Quebec, which relies predominantly on electric heating, we project declines in peak demand despite growth in summer demand. In most provinces, however, we project a switch from winter-peaking to summer-peaking electricity grids turning a notoriously cold Canada into a summer-peaking country by end-century.

A relatively unexplored area of the literature is how climate change will change the intraday shape of demand. We find a large and universal increase in ramping requirements—the range between minimum and maximum hourly demand within a day—increasing the need for greater flexibility in future electricity systems. This finding from the demand side echoes a similar need coming from the supply side, where an increasing share of variable energy renewable resources is placing a greater importance on flexibility to meet larger ramping requirements.

[1] Previous literature has explored the effect of climate change on mortality (Barreca et al., 2016; Heutel et al., 2017), economic growth (Dell et al., 2012), economic production (Burke et al., 2015) and human capital (Graff Zivin et al., 2018), to name a few. In terms of future energy demand, several studies have examined the nonlinear effect higher temperatures are expected to have on electricity demand (Auffhammer and Aroonruengsawat, 2011; Auffhammer et al., 2017; Davis and Gertler, 2015; Wenz et al., 2017). Kahn (2016) offers a review of the climate change adaptation literature.

Our empirical analysis consists of two parts. First, we estimate the relationship between temperature changes and demand for each province. We call the resulting relationships *temperature response functions*, i.e. the marginal effect of temperature on electricity demand. Second, we project changes to future electricity demand by applying our estimated temperature response functions to projections of temperature changes by mid- and end-century. We report results for both the level (annual and seasonal) and timing (peak and intraday) of demand across all provinces.

To estimate the causal relationship between temperature and electricity demand, we draw on public and private sources to construct an original dataset of hourly observations of electricity demand for every Canadian province over the period 2001–2015.[2] We find temperature response functions characterized by a familiar U-shaped relationship: at colder temperatures, rising temperature leads to decreased electricity demand; whereas at warmer temperatures, rising temperature increases electricity demand. However, these estimates represent only the short run response, i.e. the assumption that future behaviour and technology matches that of today—an unsatisfactory result for long run projections.

To thus incorporate potential adaptation, we exploit the significant heterogeneity in temperature responsiveness across provinces. These differences correspond to key observed differences in underlying ways electricity is used across provinces, i.e. differences in air conditioner and electric heat penetration, and residential share of total demand. Re-estimating temperature response functions based on these key observables, allows us to estimate future temperature responsiveness at various counterfactual levels that reflect potential adaptation. We inform our adaptation-inclusive scenarios by estimating a model of air conditioner adoption based on household-level microdata. We find by end-century, under most emission scenarios, air conditioner penetration across Canada reaches nearly 100% in most provinces. Combining our model of air conditioner adoption with the above temperature response functions delivers long run adaptation-inclusive demand projections.

Our paper contributes to a new and growing literature, building on three recent studies that explore the effect of climate change on electricity demand. First, in terms of regional heterogeneity, our paper finds similar results as Wenz et al. (2017): rising temperatures do not significantly increase electricity demand in a cold country, such as Canada. However, whereas Wenz et al. (2017) focus on regional heterogeneity driven by large climatic differences between southern and northern European countries, our paper finds differences in projected demand changes within Canada, despite relatively similar climatic conditions across provinces. Instead, we find heterogeneity driven by large variation in temperature-sensitive uses of electricity (e.g.

[2] This dataset consists of a mix of publicly and privately provided data from all ten provinces in Canada. We thank representatives from multiple balancing authorities and grid operators for their willing to provide the data. To the best of our knowledge, no other hourly multi-year panel dataset of Canadian electricity demand is available.

electric heating penetration). Our finding emphasizes the importance of understanding the underlying drivers of temperature-sensitive demand.

Second, similar to Auffhammer et al. (2017), we emphasize the importance of looking beyond average effects in difficult-to-store electricity, and thus project changes in both average and peak demand. However, we go one step further in using our hourly granularity to estimate changes in the intraday shape of demand. This aspect is of critical importance to global electricity systems already grappling with large swings in intraday supply from a growing share of renewable resources. Considerable attention has been paid to the electricity "duck curve", so-named due to the shape of intraday net demand characterized by a midday belly of low net demand when solar is generating at its fullest, followed by a steep ramp in the late afternoon having the appearance of a duck's neck (CAISO, 2016). Our results provide evidence of the need for even more flexibility to manage greater intraday variance coming from the demand side as well.[3]

Third, we develop a tractable method to incorporate adaptation into future projections of temperature-induced demand changes. Similar to Davis and Gertler (2015) we model the adoption of air conditioners in response to changes in temperature using household-level microdata, which can be used to project future air condition penetration under a warmer climate. However, whereas Davis and Gertler (2015) use this information to project future demand using a temperature response function from a different region with currently high air conditioner penetration levels, we estimate temperature response directly as a function of air conditioner penetration and other temperature-sensitive observables. This innovation allows us to use the projected air conditioner penetration levels directly, while keeping region-specific fixed effects and other characteristics, to project future demand changes with adaptation.

In a recent paper, Auffhammer (2018) exploits significant cross-sectional variation at the household level to estimate the relationship between temperature sensitivity and extant climate conditions. In doing so, this approach provides a reduced form method to incorporate adaptation by making temperature response a function of prevailing climate. This is a promising straightforward approach with the requirement of significant cross-sectional data. Our method is comparable and both papers seek the same thing: the effect of changing climate on electricity demand, incorporating elements of adaptation. Our method unpacks the relationship by decomposing the change into its components: the direct effect of temperature on demand and the indirect effect of temperature altering the stock of temperature-sensitive durables, such as air conditioners, and the corresponding effect of higher levels of air conditioner penetration on demand. Thus our method offers different insights as to the channels driving the changes.

Our paper is structured as follows. In Section 2, we motivate our empirical work with a conceptual framework that breaks down the effect of temperature on demand into (i) the direct

[3]To continue with the waterfowl analogy, we provide evidence of a "stretching of the duck's neck" coming from the demand side due to higher temperatures leading to larger intraday minimum to maximum hourly demand ranges.

effect and (ii) the indirect effect incorporating adaptation. Section 3 describes our data. In Section 4, we estimate temperature response functions, building from short-run to long-run estimates that incorporate adaptation. In Section 5, we combine our estimated temperature response functions with projected temperature changes to project changes in future electricity demand. Section 6 concludes by discussing some of the implications of our findings.

2 Conceptual Framework

To motivate the empirical analysis, we posit the following representation of electricity demand that responds to temperature and other factors:

$$y = f(T, D(T), X) \tag{1}$$

The first element is the direct effect temperature T has on electricity demand y. The second term allows temperature to affect demand indirectly via $D(T)$. We can think of D as a vector of durables whose stock is both influenced by temperature and in turn alters the temperature sensitivity of demand.[4] As a concrete example, one can imagine the stock of air conditioners in a region to be an element of D. Higher temperatures directly affect the stock of air conditioners, and in turn the higher stock increases the temperature sensitivity of demand as a result of more air conditioners turning on during heat waves.[5] Conversely, one can also imagine a different element of D that has the opposite effect. For example, the stock of energy efficiency, such as better home insulation, is likely affected by changes in temperature, and a higher stock dampens the temperature sensitivity of demand. Lastly, X captures other variables that affect demand independently of temperature.

To see how temperature changes affect demand, we differentiate Eq.1 with respect to T:

$$\frac{dy}{dT} = \underbrace{f_T}_{\substack{\text{Direct effect} \\ \text{or} \\ \text{Intensive margin}}} + \underbrace{f_D \frac{dD}{dT}}_{\substack{\text{Indirect effect} \\ \text{or} \\ \text{Extensive margin}}} \tag{2}$$

Equation 2 demonstrates the components of demand response to temperature. The first term, f_T, is the **direct effect** of changing temperature *holding the stock of D constant*, i.e. $\frac{\partial f(T,\bar{D})}{\partial T}$.[6]

[4] We use the label *durables* as we largely focus on the role of air conditioners and electric heating, however, D encompasses a broader set of factors, potentially including societal norms and behaviors that influence temperature sensitivity.

[5] More precisely, higher *expectations* of future temperatures are likely to drive decisions regarding durables stock, such as air conditioners. We ignore differences in timescale here, but address them in the empirical work that follows.

[6] We use the shorthand notation f_T to represent the partial derivative of $f(T,D)$ with respect to the first element, T. Similarly, f_D represents the partial derivative of $f(T,D)$ with respect to its second element, D.

This is what Davis and Gertler (2015) call the intensive margin, or what Burke and Emerick (2016) and Graff Zivin et al. (2018) call the short-run response.

The second term is the **indirect effect**. It is the product of temperature changing the stock of durables through $\frac{dD}{dT}$ and, in turn, the change in durables stock affecting demand through f_D. The change in stock leads Davis and Gertler (2015) to refer to this mechanism as the extensive margin. Since we can imagine the timescale at which the stock of durables changes to be long, or slow-acting, we can consider the sum of the direct and indirect effects to be the long-run response.

In our empirical strategy, we set out to estimate all three objects in Eq.2: f_T, f_D and $\frac{dD}{dT}$, but first we briefly describe the data in the following section.

3 Data

3.1 Electricity demand

The analysis is made possible due to a rich new dataset of hourly electricity demand for each of Canada's ten provinces. The dataset was constructed in part from publicly available data in provinces with competitive electricity markets, but in most cases from the collection of private data directly from the respective provincial utilities and/or balancing authorities. For each province, the data consist of a time series of hourly system-wide demand over the period 2001–2015.[7] Thus, for most provinces there are roughly 131,000 observations. The hourly demand varies significantly across provinces and by season, reflecting large population differences and seasonal electricity uses that differ by province. As this is total system-wide demand, it also represents a mix of residential, commercial and industrial demand. Summary statistics are listed in Table 1.

3.2 Temperature-sensitive demand drivers

We use data on temperature-sensitive demand drivers from two sources. First, we collect data on air conditioner penetration, electric heat penetration and residential shares of total electricity demand from Natural Resource Canada's Comprehensive Energy Use Database (CEUD). The CEUD data is an annual province-level panel from 2001 to 2015 with significant cross-sectional and temporal variation (see Figure 1). Residential air conditioner penetration has grown in all provinces in the years 2001–2015, however, there remain large differences across provinces (greater than 80% in ON vs less than 10% in NL). Residential electric heat varies across provinces, but stands out in QC—a province with large (and relatively cheap) hydro-electric resources. Residential shares are roughly bi-modal, with most provinces having roughly one-third of their

[7] For PE, NS, NL and QC, the data are only available for 2007 onwards.

Table 1: Summary Statistics

	Average Demand (aMW)		Peak Demand (MW)		Mean Temp (°C)	
	Summer	*Winter*	*Summer*	*Winter*	*Summer*	*Winter*
BC	6163	7687	9061	11039	12.8	5.3
AB	7592	8295	10441	11229	8.8	-4.3
SK	2252	2623	4654	3682	8.8	-10.3
MB	2095	2929	3464	4366	9.6	-10.2
ON	15262	17461	26854	24979	11.9	-1.0
QC	18166	25179	29411	39266	11.5	-3.7
NB	1377	1949	2543	3326	10.3	-3.6
PE	131	150	205	265	10.3	-2.6
NS	1187	1511	1806	2192	10.8	-1.6
NL	619	936	1271	1523	8.9	-1.3

Notes: Summary statistics are for 2001–2015 (2007–2015 for PE, NS, NL and QC). Summer refers to April–October, winter refers to November–March. An average MW, or "aMW", is the total MWh of seasonal demand divided by the number of hours in the season.

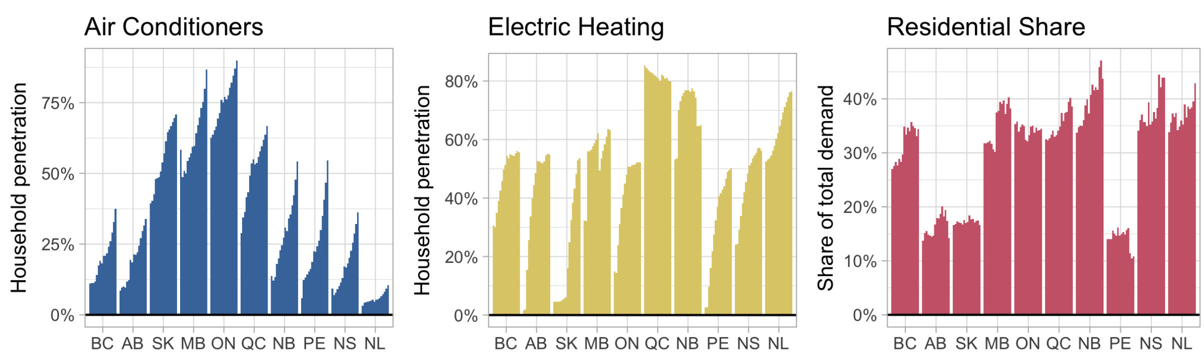

Figure 1: Observable temperature-sensitive demand drivers

Notes: Each bar represents an annual value for the years 2001–2015. Data from Natural Resource Canada's Comprehensive Energy Use Database (CEUD).

total demand attributed to the residential sector, whereas AB, SK and PE have significantly lower residential shares. In the cases of AB and SK this is due to large electricity-intensive industrial sectors, whereas PE has a large commercial sector relative to residential.

Second, for the estimation of a model of air conditioner adoption, we collect household level microdata from Statistics Canada's Household and the Environment Survey (HES). The HES data come in several waves (we use the 2006, 2007, 2009, 2011 and 2013 waves of the HES) and contain information on air conditioner ownership, income, household demographic variables, ownership or rental status, and household age and size. The data are provided at the Census Subdivision (city) level, which we can then match to temperature data at the same level.

3.3 Historical temperature

We collect data from Environment Canada to calculate population-weighted hourly temperatures for each province for each hour of the 15 year period corresponding to the demand data.We then merge the hourly temperature data with our electricity demand data. Table 1 summarizes mean seasonal temperatures across the provinces. Mean summer temperatures range from 8.8 to 12.8°C, whereas mean winter temperatures show wider variation: from -10.3° in Saskatchewan (central Canada) to +5.3° in the west coast province of British Columbia.

3.4 Projected temperature

We obtain forecasts of future temperatures based on statistically downscaled global climate model outputs from the Pacific Climate Impacts Organization at the University of Victoria.[8] For our main results, we use an ensemble of projections from 12 global climate models under from CMIP5—the Coupled Model Intercomparison Project Phase 5.[9]

The data include temperature projections for mid-century (2041-2060) and end-century (2081-2100) at a roughly 10km gridded spatial granularity. We then geo-match the individual coordinates to 2016 Canadian census population data to produce population-weighted projections at the province level. We repeat this process for two *Representative Concentration Pathway* scenarios (RCP 4.5 and 8.5), representing alternative assumptions regarding mitigation efforts.[10] RCP 8.5 represents the so-called business-as-usual scenario, where little to no mitigation of greenhouse gas emissions are taken, and correspondingly large temperature increases (roughly 5.5°C by end-century for the national average). RCP 4.5 can be considered the moderate emission scenario, with end-of-century temperature increases of roughly 3°C across Canada. For the main demand projection results we use RCP 8.5, but include results from the alternative RCP scenarios in the Appendix.

Figure 2 plots the mid- and end-century projected temperature changes from the ensemble model for both the RCP4.5 and 8.5 scenarios. The projected temperature changes differ significantly by month and province.

[8] Source: https://pacificclimate.org/data/statistically-downscaled-climate-scenarios

[9] We take the simple average of projection from the following generalized circulation models (GCMs): ACCESS1.0, CanESM2, CCSM4, CNRM-CM5, CSIRO-Mk3-6.0, GFDL-ESM2G, HadGEM2-CC, HadGEM2-LR, INM-CM4, MPI-ESM-LR, MRI-CGCM3, MIROC5. We use both the RCP8.5 (high emissions) and RCP4.5 (medium emissions) scenarios for our analysis.

[10] For a thorough overview of representative concentration pathways, see Van Vuuren et al. (2011).

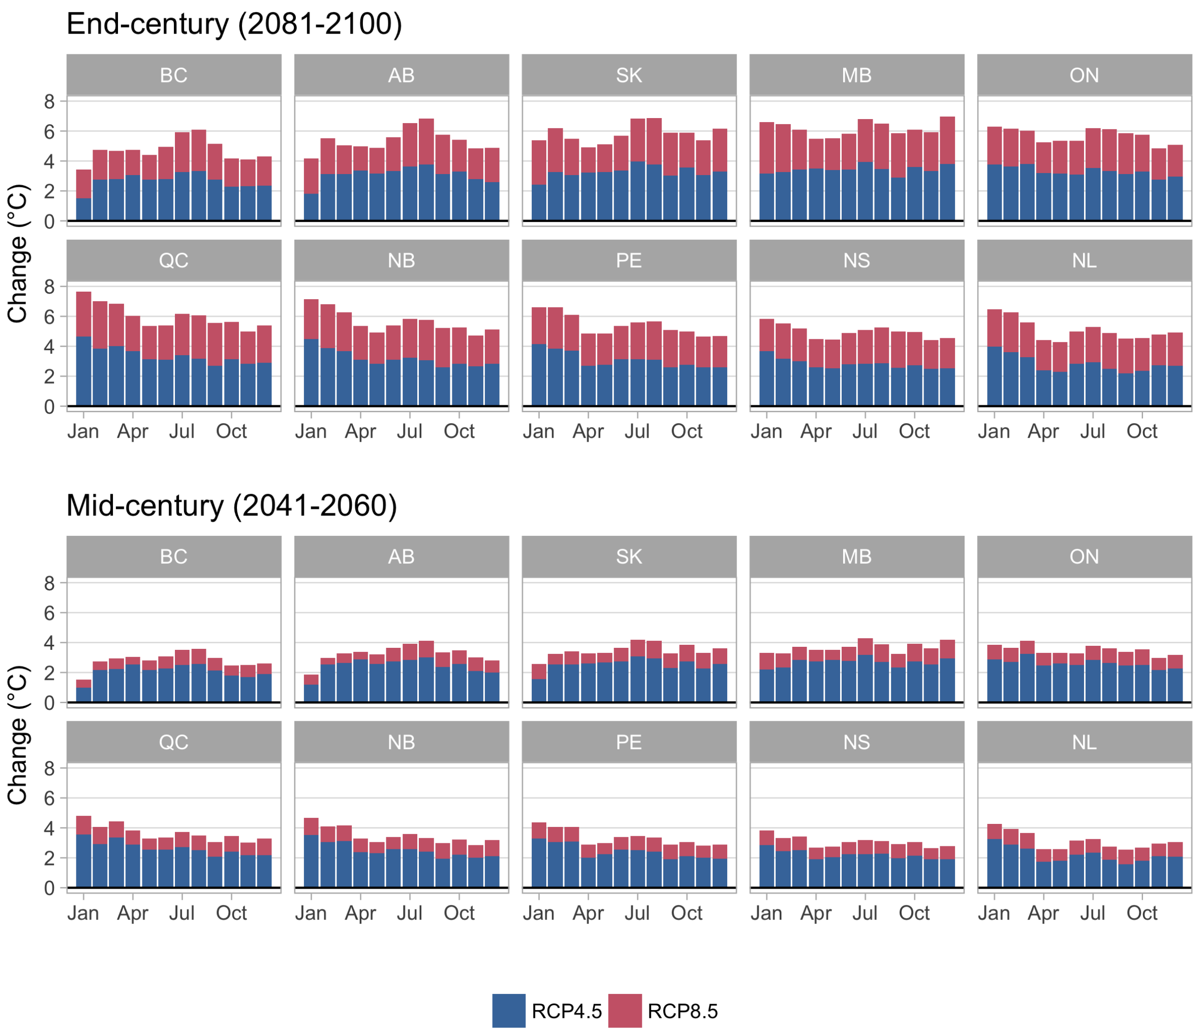

Figure 2: Projected temperature changes

Notes: Projected population-weighted temperature changes (°C) by province and month between 1981-2000 baseline and two periods: mid-century (2041-2060) and end-century (2081-2100). Based on the CMIP5 ensemble RCP8.5 (high emissions) and RCP4.5 (medium emissions) scenarios.

4 Estimating temperature response functions

Our strategy to estimate temperature response functions involves three steps. In the first step, we estimate the short run effect of temperature on electricity demand, conditional on province-specific (unobserved) drivers of demand. In step two, we re-estimate temperature response functions, this time conditioning on selected key observable drivers of demand rather than unobserved provincial differences. We demonstrate the strong fit between our estimated temperature response functions using the latter method and the province-specific method when evaluated at historical levels of the key observables. Finally, in the third step, we estimate a model of air conditioner adoption that can be used to inform the long run temperature response at higher levels of air conditioner penetration.

Step 1: Short run temperature response

We estimate the relationship between temperature and electricity demand using our hourly dataset of historical temperatures and electricity demand by province. Referring to Eq.2 of our conceptual framework, we estimate the first term f_T, exploiting the hourly data with fixed effects estimation. Specifically, we run ten separate regressions—one for each province—regressing the log of hourly electricity demand on temperature variables and a rich set of date-time fixed effects:

$$\log(y_t^p) = \sum_b \beta_b^p T_{tb}^p + \gamma^p \theta_t + \epsilon_t^p \tag{3}$$

Temperature enters semi-parametrically, with T_{tb}^p representing a dummy for whether temperature in province p at date-time t falls in the temperature bin b. Bins are defined in 2°C increments from -45°C to +39°C, the full range of hourly temperatures in Canada from 2001–2015.[11]

A large number of unobserved factors in addition to temperature influence electricity demand in any given period. In order to identify the effect of temperature on electricity demand, we employ a fixed effects strategy that controls for unobserved factors that vary predictably over time. Specifically, θ_t contains hour-of-day, day-of-week, day-of-year, statutory holiday, and year fixed effects (dummy variables). Hour-of-day dummy variables absorb systematic differences in electricity demand that occur within a day. This is important, as temperature also varies across the day. Day-of-year dummy variables soak up any variation in electricity demand that occurs over the year, such as seasonal variation in demand. Year dummy variables

[11] An alternative specification is to use the concept of heating (cooling) degree hours, which count the number of degrees below (above) an arbitrary threshold, typically 18°C. While common in the electricity literature, the semi-parametric specification allows for a more flexible non-linear response.

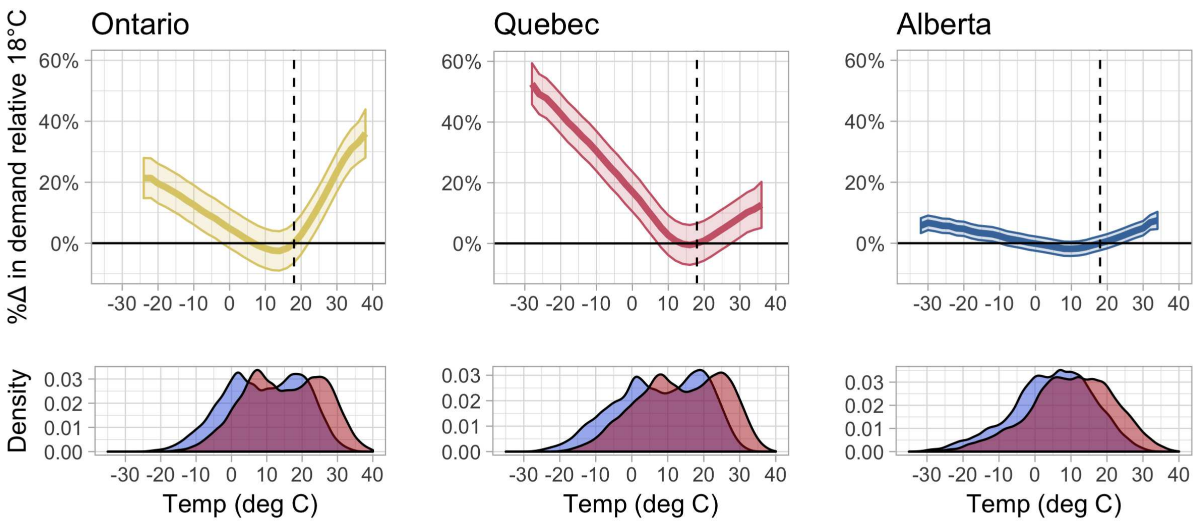

Figure 3: Temperature response functions and end-century (RCP8.5) temperature changes for 3 major provinces

Notes: The temperature response function represents the coefficients of a regression of log(demand) on 2°C temperature bins and fixed date effect controls, relative to the 17-19°C bin. In other words, the percentage change in electricity demand as temperatures differ from 17-19°C. The shaded region represents the 95% confidence interval with Newey-West heteroskedasticity and autocorrelation consistent standard errors calculated using a 168 hour (1 week) lag. The bottom panels show the density of historical (blue) and projected (red) hourly temperatures for end-century (RCP8.5).

pick up changes in demand that occur from one year to the next, for example due to changes in population or in the quality of housing stock. Day-of-week and statutory holiday dummy variables pick up variation in electricity demand that occurs across days of the week or on holiday days. Successful identification of the effect of temperature on short-run electricity demand requires that unobserved shocks to electricity demand are not correlated with temperature after conditioning on the fixed effects described above. Because of the high resolution of fixed effects covering key drivers of electricity demand that we include in our specification, as well as year fixed effects making our identification based on within-year variation, we believe that this specification should successfully identify the short-run effect of temperature on demand.[12]

We show estimation results graphically for three large provinces with distinct temperature response functions in Figure 3.[13] The interpretation of the value of the function is the percentage change in electricity demand for a given hourly temperature relative to the omitted 17-19°C bin. Ontario, with the highest share of household air conditioner penetration of all provinces in Canada, has the steepest "right-side" temperature response function slope—demand increases sharply at hot temperatures as a result of cooling demand. Whereas Quebec, with its steep

[12] Smith (2016) and Rivers (2016) use similar methods involving high resolution fixed effects to examine the causal effect of daylight saving time on traffic fatalities (Smith) and electricity demand (Rivers). Auffhammer et al. (2017) and Wenz et al. (2017) use similar methods for their identification of short run temperature response functions.

[13] Short run temperature response functions for all provinces are shown in the Appendix.

"left-side" slope is the province with the highest share of households using electric heat as their primary heating source and thus heating degree sensitivity—demand increases sharply as temperature drops *below* a nadir level of roughly 14°C. Alberta, with the highest industrial share of electricity demand of any province (and thus low share of residential demand), has a rather flat temperature response function, reflecting weak temperature sensitivity. The bottom panels in Figure 3 show historical and projected future hourly temperatures (RCP8.5 scenario) for each region at end-century.

Step 2: Re-estimating temperature response based on observables

The prior section used high resolution fixed effects estimated separately by province to allow us to cleanly identify the causal effect of changes in temperature on short-run electricity demand. However, while the fixed effects are useful for identification, they also prevent us from understanding structural reasons why temperature responses might differ. In this section, we remove the province-specific estimation, and instead estimate a single temperature response equation based on key observable temperature-sensitive characteristics that differ across space (province) and time (year).[14] This allows us to understand what drives differences in temperature responsiveness across provinces and time. Accordingly, this allows for flexible counterfactual scenarios that include adaptive behaviour in response to higher temperatures, such as increased air conditioner penetration, that in turn affect temperature sensitivity in the long run.

To clarify this approach, we show how our empirical strategy aligns with our conceptual framework with the following illustration. Consider a version of Eq.3, but as a single regression, rather than ten separate provincial regressions, and including key observables D and their interaction with temperature:

$$\log(y_{tp}) = \sum_b \beta_1^b T_{tpb} + \beta_2 D_{tp} + \sum_b \beta_3^b T_{tpb} D_{tp} + \gamma\theta_t + \eta_p + \epsilon_t \tag{4}$$

For the sake of building intuition, we simplify Eq.4 by dropping the semi-parametric binned temperature notation and revert to a generic T purely for exposition:

$$\log(y_{tp}) = \beta_1 T_{tp} + \beta_2 D_{tp} + \beta_3 T_{tp} D_{tp} + \gamma\theta_t + \eta_p + \epsilon_t \tag{5}$$

Differentiating Eq.5 with respect to T gives the marginal effect of temperature on demand. Rearranging highlights the equivalency of this empirically-estimable equation to our conceptual

[14] Another way to think of our previous estimation method is as a single regression whereby province dummies are interacted with temperature bins and fixed effects. In this section's specification, we instead interact temperature variables with the vector of observables, thus any differences in temperature responsiveness are explained by differences in observables rather than unobserved provincial heterogeneity.

framework:

$$\begin{aligned}\frac{\partial \log y}{\partial T} &= \beta_1 + \beta_2 \frac{\partial D}{\partial T} + \beta_3 D(T) + \beta_3 T \frac{\partial D}{\partial T} \\ &= \underbrace{\beta_1 + \beta_3 D(T)}_{f_T(T,D)} + \underbrace{(\beta_2 + \beta_3 T)}_{f_D(T,D)} \frac{\partial D}{\partial T}\end{aligned} \tag{6}$$

Thus, our challenge is to estimate the above β's to estimate temperature response for a given level of observable characteristics (D). To do so, we regress electricity demand on temperature and observables as per Eq.5, replacing the generic temperature notation with heating and cooling degree variables:

$$\log(y_{tp}) = \beta_{11} CD_{tp} + \beta_{12} HD_{tp} + \beta_2 D_{tp} + \beta_{31} CD_{tp} D_{tp} + \beta_{32} HD_{tp} D_{tp} + \gamma\theta_t + \eta_p + \epsilon_t \tag{7}$$

where CD_{tp} and HD_{tp} are cooling and heating degrees, i.e. the number of degrees actual temperature is above and below, respectively, a neutral temperature baseline.[15] This is a slightly less flexible specification than temperature bins, but still allows for different trends on either side of the neutral temperature baseline. It also greatly simplifies the regression and delivers easily interpretable coefficients.

This regression specification requires taking a stand on the elements of the temperature-sensitive durables vector, D. We include air conditioner and electric heating penetration levels and residential share of electricity demand, which we observe for each province-year. We again control for time fixed effects as well as province fixed effects that are, importantly, no longer interacted with temperature. Thus, the heterogeneous effect of temperature on demand across provinces comes only through differences in the observable characterized represented by D.

We estimate multiple variants of Eq.7, with results listed in Table 2. The first specification is the most straightforward: we include cooling and heating degrees, elements of D (air conditioner penetration, electric heating penetration and residential share) and the interaction between the temperature variables and observables. The second model augments the first by interacting residential share with air conditioner and electric heat penetration. This allows the potential for greater effect of durables at higher shares of temperature-sensitive residential demand. Newey-West heteroskedasticity and auto-correlation consistent standard errors are calculated with a 168 hour (one week) lag, as per the short run estimates.

Looking at Column 1, the sign of the coefficients is as expected: the sensitivity of demand to cooling degrees (temperature above 14°C) increases with greater air conditioner penetration (as seen by the significantly positive coefficient on $cd \times AC$). Similarly, the sensitivity of demand to

[15]We use 14°C as our neutral temperature baseline as this is the observed average nadir of the previously estimated short-run temperature response functions across Canada.

Table 2: Regression estimates of demand on observables

	Dependent variable:	
	log(load)	
	(1)	(2)
Cooling degrees (cd)	0.007*** (0.002)	0.018*** (0.003)
Heating degrees (hd)	−0.008*** (0.001)	−0.009*** (0.001)
cd×AC	0.026*** (0.001)	−0.002 (0.007)
cd×Res Share	−0.031*** (0.006)	−0.066*** (0.011)
cd×AC×Res Share		0.091*** (0.024)
hd×Electric Heat	0.021*** (0.001)	0.021*** (0.003)
hd×Res Share	0.024*** (0.002)	0.025*** (0.008)
hd×Electric Heat×Res Share		−0.001 (0.012)
Observations	1,214,730	1,214,730
Adjusted R^2	0.996	0.996

Notes: *p<0.1; **p<0.05; ***p<0.01

Newey-West HAC standard errors in parentheses.

heating degrees (temperature below 14°C) increases with greater electric heating penetration (as seen by the coefficient on $hd \times ElectricHeat$). In Column 2, the interaction term between air conditioner penetration and cooling degrees is rendered insignificant, but the effect is observed via the strongly positive triple interaction term with residential share. Cooling demand is increasingly sensitive to air conditioner penetration at higher levels of residential share.

Using these estimates, we can calculate temperature response functions holding the stock of durables constant at historical averages. Figure 4 plots the predicted temperature response functions for three provinces, holding the elements of D at historical average levels, as compared to the short run temperature response functions estimated separately by province. This figure highlights the strong explanatory power of this rather small set of observables in explaining the heterogeneity across province-specific temperature sensitivity.

A key concern with this method is the possibility of omitted variable bias.[16] We cannot exclude this possibility entirely, however, we address this issue in several ways. First, we demonstrate the strong fit of the predicted temperature response functions at historical observables levels with the province-specific short run temperature response functions previously estimated. This is shown graphically in Figure 4.

[16] The problem of *selection on observables* is not uncommon in empirical research. Oster (2017) proposes a method to investigate the likelihood of bias due to selection on observables by generalizing the approach previously suggested by Altonji et al. (2005). This involves estimating a *coefficient of proportionality*, δ, to determine how explanatory the unobservables would have to be to render the coefficients of interest insignificant. However, this approach is not entirely appropriate for our context. In the Oster (2017) *selection on observables* problem, the concern is whether controls have been appropriately selected such that the coefficient of interest on the treatment variable is robustly estimated. In our case, the selected observables are themselves the variables of interest, not simply controls.

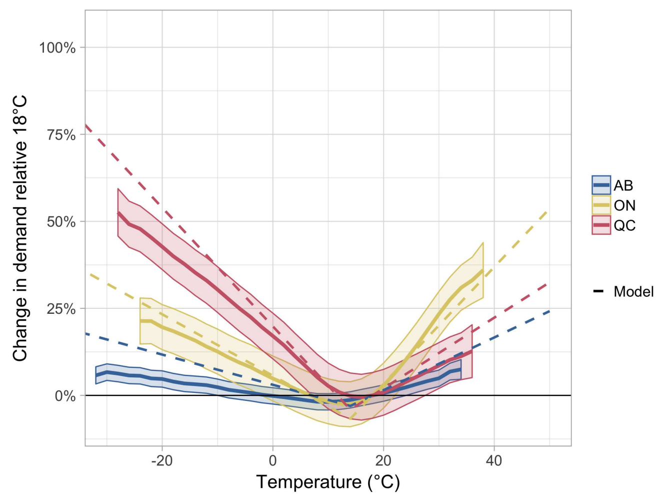

Figure 4: Comparison of short run and predicted temperature response functions evaluated at historical averages

Notes: Temperature response functions show the estimated effect of temperature on electricity demand relative to 18°C. The short run estimates are shown by the solid lines with 95% confidence intervals (shaded). The predicted temperature response functions evaluated at historical average levels of air conditioner penetration, electric heating penetration and residential share of demand are shown by dashed lines (confidence intervals omitted on the latter for figure clarity).

Second, we take a statistical approach. Our goal is to determine whether the selected observables in D sufficiently explain the underlying unobservable heterogeneity driving differences in short run temperature response functions when estimated separately by province. Thus, we take a straightforward approach to determining the appropriateness of the selected observables by regressing the slopes of the short run temperature response functions (a linear proxy for the estimated sensitivity to heating and cooling degrees) on the selected observables and examining the degree to which the fitted results explain the variation in the data, i.e. the R-squared.

The explanatory power of this small set of observables—air conditioner penetration, electric heat penetration and residential share of demand—is strong, showing an R-squared of 77-85% depending on specification. Despite only 10 observations (one for each province) the adjusted R-squared remains relatively strong despite only three observable variables used to explain the provincial heterogeneity. The full goodness-of-fit regression results and added-variable residual plots are shown in the Appendix A.3.

Step 3: Modelling air conditioner adoption

The last part of obtaining the long run response involves estimating $\frac{\partial D}{\partial T}$—the change in durables in response to changing temperature. We focus solely on the effect of higher temperatures on air conditioner penetration; we do not estimate temperature driven changes to electric heating penetration or residential share of electricity since we consider these variables to be largely driven by policy and economic factors rather than temperature.[17]

We estimate a model of air conditioner adoption using household level microdata from Statistics Canada's Household and the Environment Survey (HES) on air conditioner penetration, in a similar approach as Davis and Gertler (2015). Specifically, we use cross-sectional variation in temperature to identify the effect of climate variables on air conditioner penetration, while conditioning on other variables. We use several waves of the HES public use microdata files, and extract data on air conditioner ownership, income, household demographic variables, and household size.[18] We obtain the Census Subdivision (city) for each household in the survey, and use historical weather data from Environment Canada to obtain measures of the climate in each Census Subdivision.

$$AC_{ict} = \delta_0 + \delta_1 \tilde{T}_c + H_i\theta + \psi_t + \nu_{ict} \tag{8}$$

where AC_{ict} is a binary variable that takes on a value of one if the household owns an air conditioner and zero otherwise, H_i is a vector of observed household covariates,[19] and ψ_t is a time fixed effect to account for changes in household air conditioner penetration over time that are common across regions. The variable $\tilde{T}_c$ captures the exposure of city c to hot temperatures. We measure the climate of cities using several different variables: (1) the highest monthly mean temperature observed between 2000 and 2005, (2) the highest daily maximum temperature observed between 2000 and 2005, (3) the mean temperature in July and August observed between 2000 and 2005, and (4) the average of the maximum daily July and August temperature observed between 2000 and 2005.[20] We estimate the model using both linear probability, with and without sampling weights provided by Statistics Canada, as well as probit models. We also estimate a model that includes province fixed effects, such that the identification of the effect of climate on air conditioner penetration is identified on within-province variation in climate. This helps to purge the data of any province-specific factors (e.g., regulations, norms) that drive air conditioner penetration.

We highlight the empirical relationship between residential air conditioner penetration and climate in Figure 5. The top panel summarizes air conditioner penetration in each of the 33 cities contained in the Households and the Environment Public Use file, as well as the average

[17] We include, however, the mechanical effect that increased air conditioner penetration would have on residential share, all else equal, and modify residential share accordingly. Specifically, the modified residential share is equal to the old residential share * $(1 - AC_{old} * Avg\_AC\_per\_HH)/(1 - AC_{new} * Avg\_AC\_per\_HH)$.

[18] We use the 2006, 2007, 2009, 2011, and 2013 waves of the HES.

[19] Household covariates include a binary variable that indicates whether the house is owned or rented, a categorical variable that captures the level of education, a variable that captures the dwelling type (apartment or home) of the dwelling, a variable that captures the number of people living in the household, indicator variables for the presence of individuals aged 0-17, 18-64, and 65-plus in the household, and a categorical variable capturing household income.

[20] We focus on the years 2000 to 2005 because we were able to assemble a complete set of weather observations over this period for all cities in our sample, without any entry or exit of weather monitoring stations.

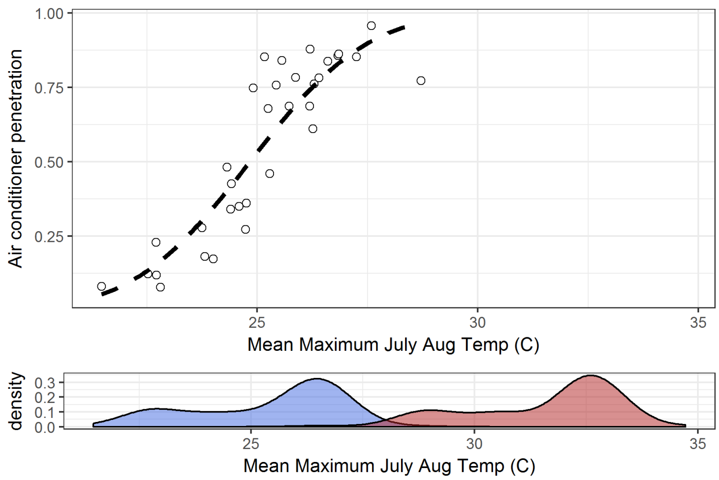

Figure 5: Air conditioner penetration as a function of climate

Notes: Top panel shows air conditioner penetration for a cross-section of individual census metropolitan areas (weighted average of the 2006-2013 HES waves) plotted against a measure of hot summer temperature (average maximum July-August temperature for 2000-2005). Bottom panel shows the distribution of mean maximum July-August temperatures historically (blue) and projected end-century in the RCP8.5 scenario (red).

daily maximum July-August temperature observed between 2000 and 2005 in each city. There is a clear positive relationship between these two variables, which is summarized by a probit fit without any covariates (regression results for probit and OLS estimates including covariates are listed in the Appendix). The bottom panel shows the current exposure to hot summer weather weighted by population, as well as the projected exposure to different climates at the end-of-century under an RCP8.5 scenario. The cross-sectional relationship between current climate and air conditioner penetration suggests that future warming will induce substantial increases in air conditioner penetration.

Motivated by the relationship in Figure 5, we estimate the relationship between climate and air conditioner penetration using the household-level data as described above (results shown in the Appendix). We find that for each 1 degree Celcius increase in the maximum daily July-August temperature, the penetration of residential air conditioners increases by 16.8 percentage points. Using alternative definition for the hot temperature variable $\tilde{T}$ also delivers positive and highly statistically significant relationship between the prevalence of warm weather and the penetration of air conditioners. We include several robustness checks as well as a probit model specification in the Appendix. We use our air conditioner adoption model to estimate projected air conditioner ownership under different future climates. Under the RCP8.5 scenario, we estimate an air conditioner penetration that increases nation-wide from about 55% today to above 99% in 2100.

Putting it all together: Long run temperature response functions

The above method of estimating temperature response response functions based on observable characteristics allows us to run scenarios of future adaptation, whereby higher levels of air conditioner penetration increase temperature sensitivity to warmer temperatures. This increased sensitivity is reflected in steepening right-side slopes of the temperature response functions. Figure 6 plots how temperature response functions for AB, ON and QC change as air conditioner penetrations increase from historical levels towards 100%.

Our multi-part method to incorporate adaptation extends the work by Davis and Gertler (2015) who estimate $\frac{\partial D}{\partial T}$ for air conditioners in Mexico. In that case, future air conditioner penetration is estimated based on projected temperature and income changes, however, the effect of changing durables on demand is not estimated. Instead, Davis and Gertler (2015) apply temperature response functions from a region with currently high air conditioner penetrations, similar to their future projections. By modelling temperature response functions as functions of the observable characteristics themselves rather than having to rely on using a comparable region's temperature response, our method allows for greater scenario analysis flexibility and the ability to retain unobservable characteristics of each province.

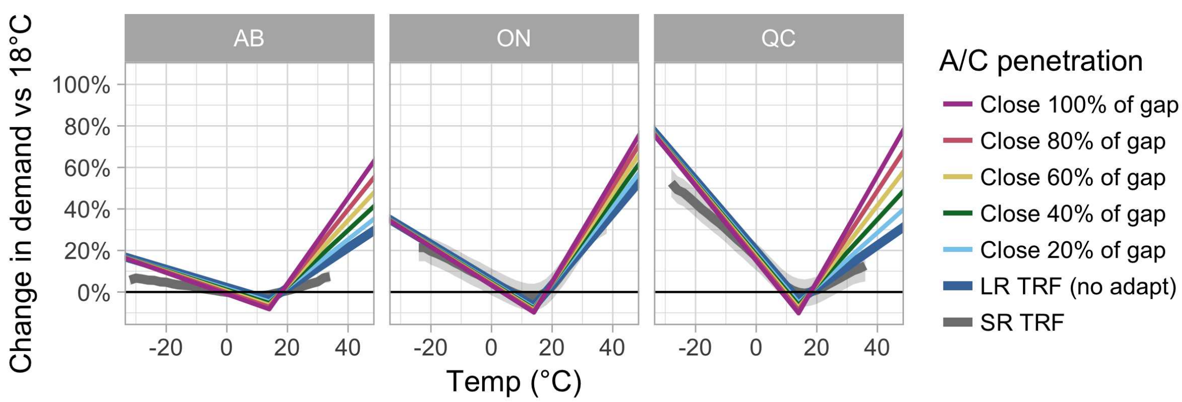

Figure 6: Temperature sensitivity at various levels of air conditioner penetration

Notes: The thick grey lines represent temperature response functions estimated separately by province. Thick blue lines represent temperature response functions estimated conditional on observables, rather than province-specific, evaluated at historic average levels. Their overlap highlights the strong explanatory power of these 3 key observables. The remaining lines represent counterfactual temperature response functions at increasing levels of air conditioner penetration, whereby the gap closes between historical and full (100%) penetration.

5 Projecting future demand changes

This section combines the temperature response functions developed in the preceding section with climate model temperature projections to project temperature-driven changes to future electricity demand.[21] We present results both without adaptation (based on the short run temperature response estimates) and with adaptation based on our estimate of increased air conditioner penetration. The main results are for the end of century under the RCP8.5 emissions scenario, with projections using alternative RCP scenarios and for mid-century available in the Appendix.

Previous literature estimating the effect of climate change on energy demand has focussed on projections of annual or seasonal demand (De Cian and Wing, 2017; Davis and Gertler, 2015). Given the instantaneous nature of electricity, and relative lack of storability, considering peak demand is also important. Accordingly, Auffhammer et al. (2017) and Wenz et al. (2017) consider the implications of climate change on both aggregate and peak electricity demand. We go one step further, exploring how the intraday shape of electricity demand will change in the future. Intraday shape is important due to, again, electricity's relative lack of storability. A steeper "ramp", i.e. the change in demand from the lowest demand hour to the highest within a day, imposes higher system costs, requiring more flexibility to manage.

5.1 Average demand

We start by presenting our results for Canada as whole, looking at monthly percentage changes in electricity demand arising under both with and without adaptation in the form of more air conditioner penetration. Looking at monthly changes highlights the effect of warming winters on reducing demand, as well as the amplified effect in the summer months coming from increased air conditioner adoption. Figure 7 plots percentage changes by month for Canada as a whole, for both short run (without adaptation) and long run (with adaptation) responsiveness, for the high emissions scenario at end-century.

Taking a closer look at individual provinces, Figure 8 shows projected changes to annual and seasonal average demand—applying both short run and long run temperature responses—for the high emissions scenario at end-century.[22] In the summer months, average demand increases across all provinces since the range of daily temperatures in these months is largely warm

[21] To project demand based on out-of-sample temperature projections (i.e. higher than previously observed), we include a linear trend term above 18°C in the specifications using temperature bins. We choose 18° by visual inspection based on where a clear linear trend is established, slightly to the right of the low demand nadir of 14°. Robustness checks to alternative thresholds, and even the inclusion of a multi-point spline, do not alter the demand projections significantly. The temperature response functions conditioned on cooling degrees do not require this modification.

[22] Detailed results, as well as those for mid-century and the medium emissions scenario (RCP4.5) are listed in the Appendix (Tables A1 & A2).

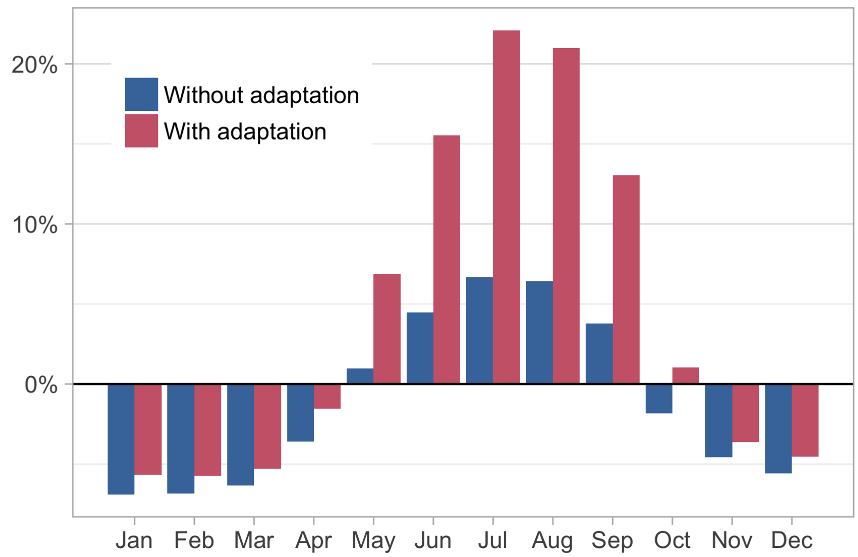

Figure 7: Monthly average demand change (RCP8.5, End-century)

Notes: Percentage change in monthly demand at end-century under the high emission (RCP8.5) scenario. The left (blue) bar uses the short run temperature response. The right (red) bar incorporates adaptation from increased air conditioner adoption to provide a long run response. Note that the level of demand differs across the months of the year (summer is smaller than winter). Thus summing the net of the bars overstates the annual percentage change. The annual change in this scenario is 3.6% (see Appendix Table A1).

enough to be located on the upward sloping portion of the temperature response functions where higher temperature *increases* electricity demand. The effect is largest in Ontario, where both the sensitivity to cooling degrees is steepest and the starting summer temperature levels are among the highest of any province. British Columbia, despite its flatter sensitivity to cooling degrees (as compared to Ontario), also shows a significant increase. This is due to its warmer average climate, meaning fewer hours in the domain of heating demand, where higher temperature decreases demand.

In the winter months, average demand declines across Canada. The effect is largest in Quebec, New Brunswick and Newfoundland & Labrador, the three provinces with most electric heating share and correspondingly steepest "left side" slope of their temperature response functions.

For the country as a whole, the estimated change in annual electricity demand is small and stands in contrast to previous studies in warmer countries showing large increases. Using the short run temperature response, annual demand falls by 1.4%. When adaptation is incorporated, whereby air conditioner penetration reaches nearly 100% nationally, annual Canadian electricity demand still only increases by 3%. This result speaks to the beneficial effect (from an energy use perspective) of a warmer winter reducing heating demand and nearly offsetting the entirety of the incremental summer cooling demand.

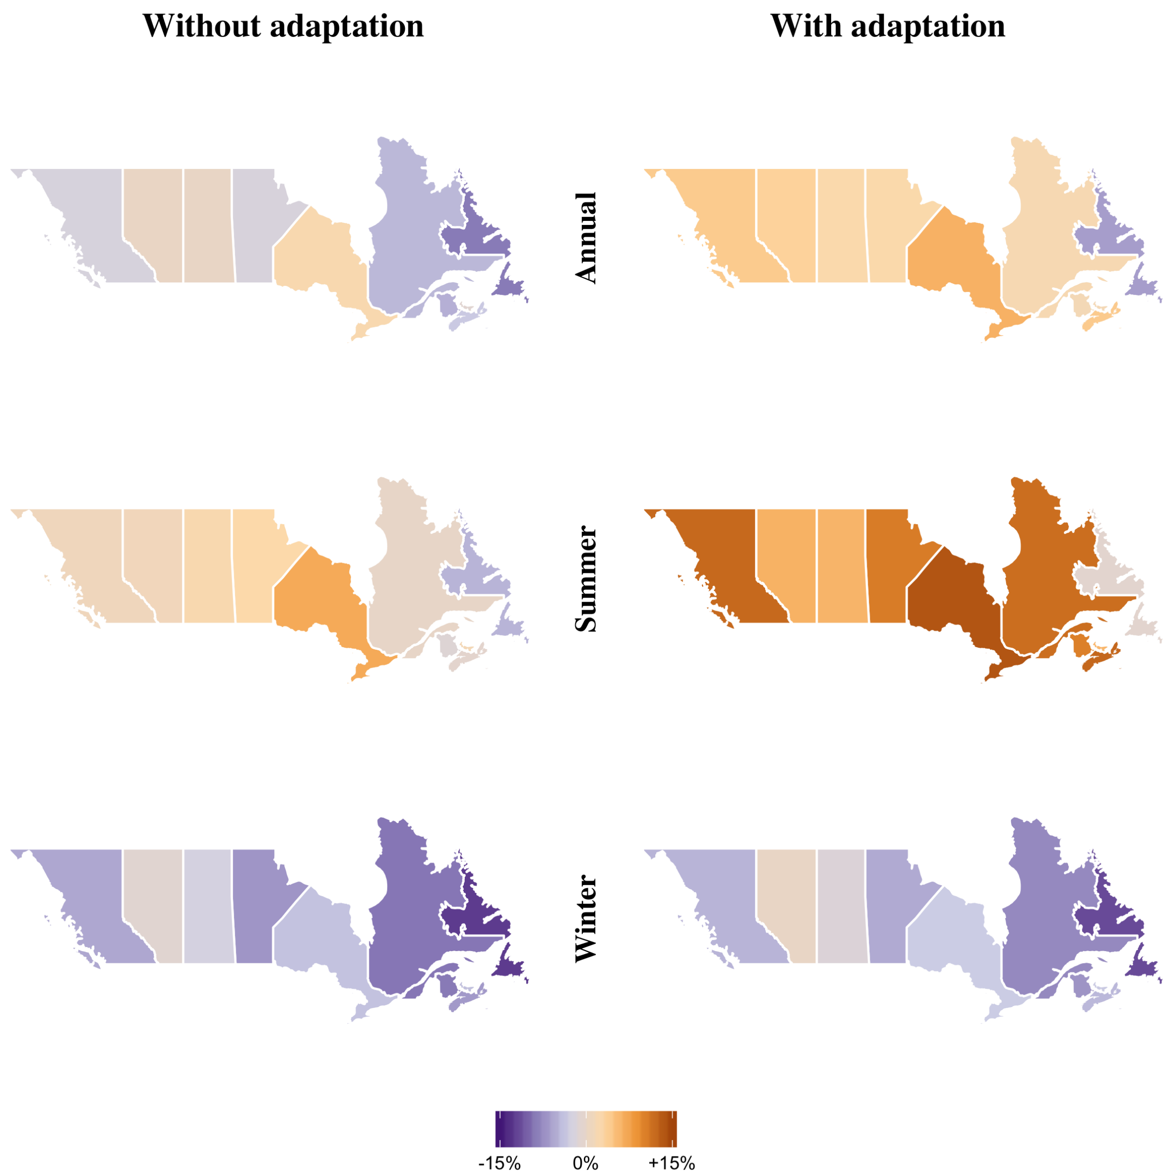

Figure 8: Annual and seasonal demand change (RCP8.5, End-century)

Notes: Maps show the percentage change in annual and seasonal total demand by province at end-century under the high emission (RCP8.5) scenario. The left column uses the short run temperature response. The right column incorporates adaptation from increased air conditioner adoption to provide a long run response.

5.2 Peak demand

The importance of peak demand relates to the generation capacity requirements of the electricity system. Lack of storability means that the system must have sufficient capacity (or capability to generate) during peak demand periods. The impact on the electricity system therefore depends not simply on how much demand increases, but when. An increase during non-peak period has no effect on peak capacity requirements.

Most provinces in Canada are winter peaking. As such, rising temperature reduces peak capacity demands in most provinces using only short run responses (see Figure 9 and Tables A3 & A4). Peak demand increases, however, for the two summer-peaking provinces of Saskatchewan and Ontario.

When we incorporate adaptation in the form of greater air conditioner penetration, the increase in summer-peaking provinces is amplified. Ontario's peak demand increase is not 35% (vs 10% using only short run responsiveness) implying roughly 10GW of new needed generating capacity in that province solely due to temperature. Furthermore, many provinces see their seasonal peak flipping from winter to summer leading to peak demand increases. BC, AB, MB and NS all go from winter peaking to summer peaking electricity systems (Figure 10). Alberta, in particular, sees a large increase (15%) in peak demand due to a significant increase in air conditioner adoption and relatively high hourly temperatures in the peak of summer.

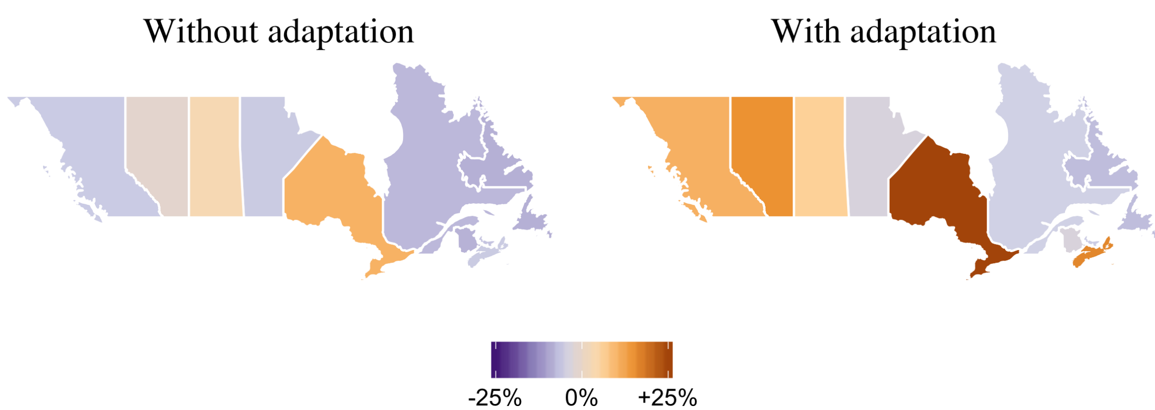

Figure 9: Peak demand change (RCP8.5, End-century)

Notes: Maps show the percentage change in peak hour demand by province at end-century under the high emission (RCP8.5) scenario. The left column uses the short run temperature response. The right column incorporates adaptation from increased air conditioner adoption to provide a long run response.

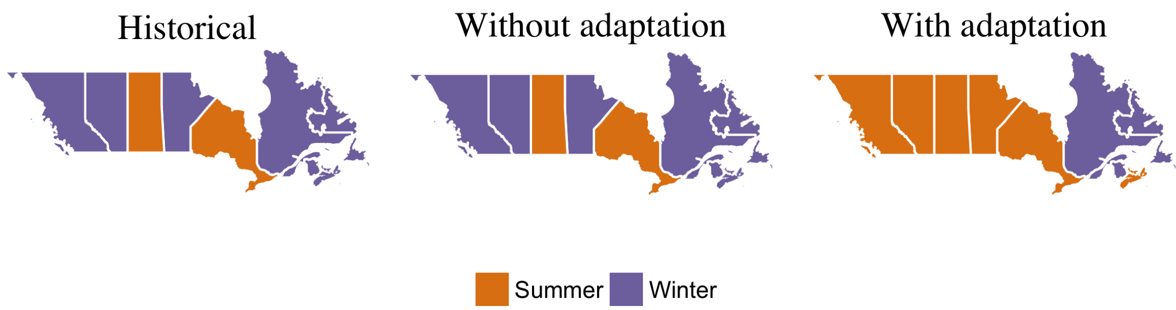

Figure 10: Seasonal peak hour demand (RCP8.5, End-century)

Notes: Maps show the period during which the annual hourly peak occurs, by province at end-century under the high emission (RCP8.5) scenario. The first column shows the historical peak, with only two provinces (SK and ON) having summer peaks. The second column uses the short run temperature response. The third column incorporates adaptation from increased air conditioner adoption to provide a long run response.

5.3 Intraday demand

Having considered the effect on average and peak annual demand, we investigate the effect of rising temperatures on the intraday shape of demand. Of note, for this analysis we use only uniform (projected) changes in temperature across the day. The non-uniform intraday effects are due solely to the non-linear temperature response functions.[23]

Figure 11 shows the change in the intraday demand profile for the province of Ontario during the summer. This figure shows the average hourly demand over the summer period for each hour of the day. There is a clear lift in the level of demand, with demand increasing across all hours, but the increase is clearly larger in the peak demand hours of the afternoon as compared to the morning hours. As a result, there is a significant increase in the intraday "ramp", i.e. the difference between the minimum and maximum demand within a day. This finding is exacerbated with long run response projections that incorporate higher levels of air conditioner penetration.

To provide summary statistics for each province, Figure 12 plots the change in "min-to-max" ramp requirements, in percentage terms, for all provinces for each month of the year. Unlike the effect on average or peak demand, the effect on intraday ramp requirements is consistent across the provinces: all provinces show an increase in min-to-max range in the summer months. In winter, most provinces see a slight decrease in ramp requirements with the intraday shape of demand getting flatter. Using short run responses, the larger increases are in the shoulder months, namely May and October. This is due to lower temperatures during the morning

[23]Peak demand occurs during the afternoon when temperatures are typically higher than off peak early morning hours. As a consequence, peak hours are more likely to be located on the upward sloping portion of the temperature response functions—in the domain of cooling demand—where rising temperatures increase demand.

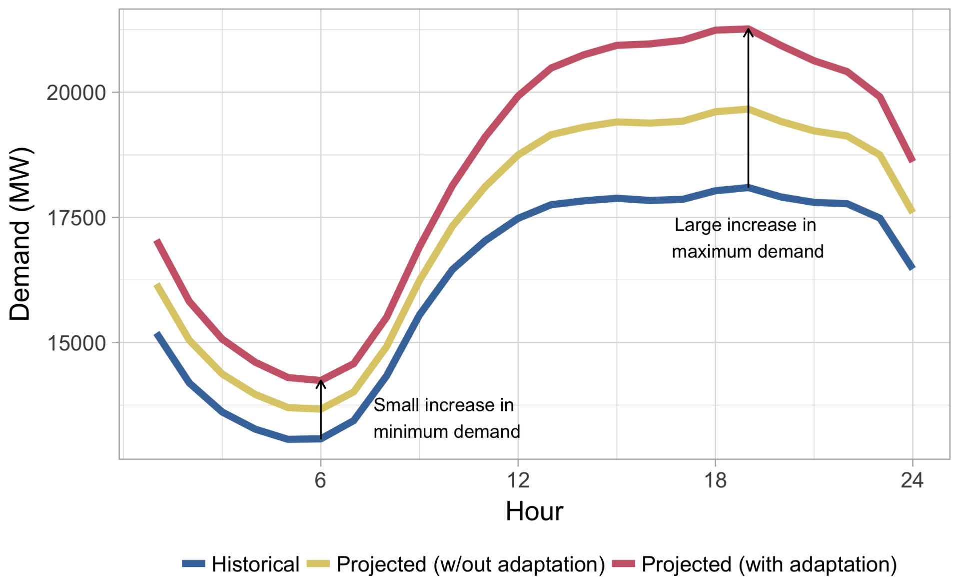

Figure 11: Intraday shape of summer demand in Ontario, historical and projected (End-century, RCP8.5)

Notes: Historical and projected end-of-century intraday demand using RCP8.5 temperature projections and authors' estimated temperature response functions, with and without adaptation.

hours falling near the low point of the temperature response functions leading to little effect on demand from rising temperatures, whereas the peak hours in those months are of sufficiently high temperature to be on the upward sloping part of the temperature response functions. Using long run responsiveness, the increase is not confined to the shoulder months, as the steeper slopes of the temperature response functions lead to significant peak demand increases in the summer. For some provinces ramping requirements increase by as much as 100%, i.e. a doubling.

This finding, coming from the demand side, adds to the projected need for more flexibility on electricity grids coming from changes on the supply side of the market, where the cost of variable renewable energy is falling and their share is growing. The so-called "duck curve" in California summarizes this issue: more solar generation in the middle of the day leads to a steep ramp in net demand (i.e. actual demand net of renewable generation) in the afternoon (CAISO, 2016). On the demand side, recent analysis from the California Energy Commission (CEC, 2018) has shown the electric vehicle charging is expected to be concentrated at residences when drivers return home from work, exacerbating the problem of meeting net demand as solar fades in the afternoon. Our finding highlights another potential issue: higher temperatures increasing ramping requirements, exacerbating both the EV charging and duck curve issues. In effect, we find higher temperatures have the potential to "stretch the duck's neck".

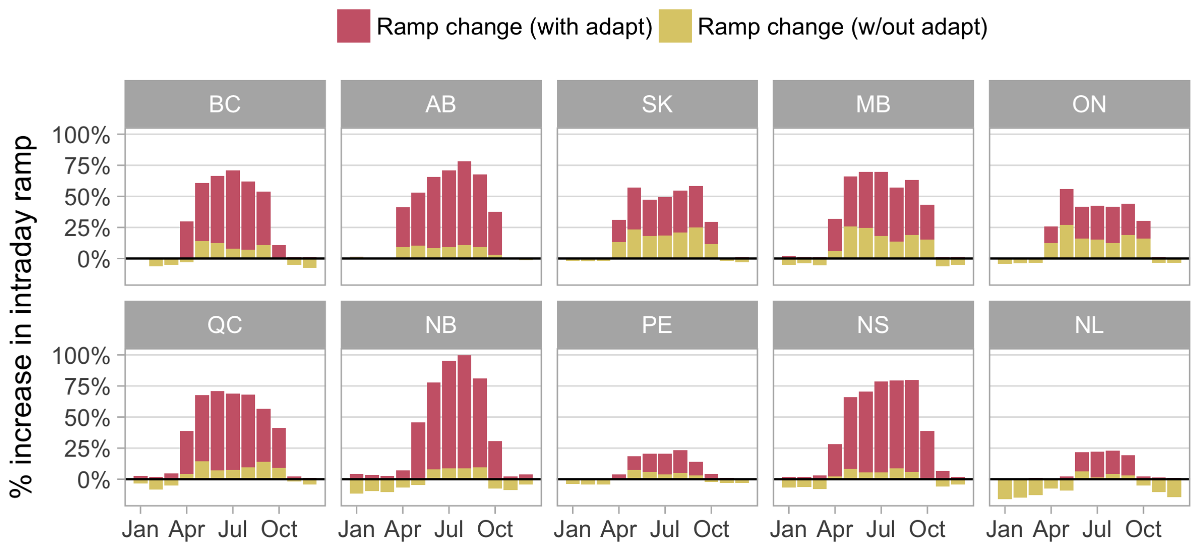

Figure 12: Percentage change in min-to-max intraday demand (End-century, RCP8.5)

Notes: Plots show the percentage change in minimum to maximum intraday demand as a result of rising temperatures across all Canadian provinces, with and without adaptation. Forecast period 2081-2100. RCP8.5 (high emissions) scenario.

6 Conclusion

This paper finds that for a colder country, such as Canada, the effect of rising temperatures due to climate change is unlikely to result in large increases in overall electricity use. In the absence of adaptation, we find only a 0.8% increase in national electricity demand in the high emissions (RCP8.5) scenario. Incorporating an increase in air conditioner adoption as a result of higher temperatures, this increase to 3.6%. Given Canada's colder climate, it is perhaps not surprising to see a much smaller effect as compared to other results in the literature that have focussed on warmer climates.

In terms of peak demand, the results are mixed across the provinces. Ontario, a summer-peaking province, sees peak demand increase by 10% in the *without adaptation* scenario and 35% in the *with adaptation* scenario. Provinces with significant electric heating, such as Quebec, see peak demand decline even in the *with adaptation* scenario. Whereas in several provinces, the annual peak demand switches from winter-peaking to summer-peaking.

It is ambitious to project estimated costs as a result of end-of-century demand changes, but as a rough estimate we can use current values for peaking capacity. At $1,000 per kilowatt, the aggregate increase in peak demand across Canada would require an investment of roughly $13 billion (USD).

An important aspect of projected demand changes arising from higher temperatures is the effect on intraday demand. We find that "ramping" requirements—the ability to swing from low to high demand within a day—is expected to increase substantially across all provinces. This result adds to parallel concerns over the need for more flexibility coming from the demand side. While we do not place a cost estimate on this effect, it speaks to the increasing value of flexibility—be it in the form of storage, peaking capacity or load-shifting—to better manage an increasing variable supply and wider-ranging demand on future electricity systems.

In sum, our paper adds to the growing literature quantifying the effects of higher temperatures arising from climate change on important economic variables, in this case electricity demand. We provide a method to incorporate adaptation by estimating temperature response functions as functions of key temperature-sensitive observables coupled with a model of air conditioner adoption at the household level. Our finding regarding intraday demand emphasizes the value and importance of capacity and flexibility, as well as the importance of understanding more than average effects when it comes to difficult-to-store electricity.

References

Akpinar-Ferrand, E. and Singh, A. (2010). Modeling increased demand of energy for air conditioners and consequent co2 emissions to minimize health risks due to climate change in india. *Environmental Science and Policy*, 13(8):702 – 712.

Altonji, J., Elder, T., and Taber, C. (2005). Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of Political Economy*, 113(1):151–184.

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. National Bureau of Economic Research Working Paper 24397.

Auffhammer, M. and Aroonruengsawat, A. (2011). Simulating the impacts of climate change, prices and population on California's residential electricity consumption. *Climatic change*, 109(1):191–210.

Auffhammer, M., Baylis, P., and Hausman, C. H. (2017). 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):1886–1891.

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., and Shapiro, J. S. (2016). Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century. *Journal of Political Economy*, 124(1):105–159.

Burke, M. and Emerick, K. (2016). Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy*, 8(3):106–40.

Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239.

CAISO (2016). What the duck curve tells us about managing a green grid.

CEC (2018). Staff report - California plug-in electric vehicle infrastructure projections. Technical report, California Energy Commission.

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*, 112(19):5962–5967.

De Cian, E. and Wing, I. S. (2017). Global energy consumption in a warming climate. *Environmental and Resource Economics*.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Graff Zivin, J., Hsiang, S., and Neidell, M. (2018). Temperature and human capital in the short and long run. *Journal of the Association of Environmental and Resource Economists*, 5(1):77–105.

Heutel, G., Miller, N. H., and Molitor, D. (2017). Adaptation and the mortality effects of temperature across US climate regions. National Bureau of Economic Research Working Paper 23271.

Isaac, M. and van Vuuren, D. P. (2009). Modeling global residential sector energy demand for heating and air conditioning in the context of climate change. *Energy Policy*, 37(2):507 – 521.

Kahn, M. E. (2016). The climate change adaptation literature. *Review of Environmental Economics and Policy*, 10(1):166–178.

Oster, E. (2017). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business and Economic Statistics*, pages 1–18.

Rivers, N. (2016). Does daylight savings time save energy? Evidence from Ontario. *Environmental and Resource Economics*, pages 1–27.

Smith, A. C. (2016). Spring forward at your own risk: Daylight saving time and fatal vehicle crashes. *American Economic Journal: Applied Economics*, 8(2):65–91.

Van Vuuren, D. P., Edmonds, J., Kainuma, M., Riahi, K., Thomson, A., Hibbard, K., Hurtt, G. C., Kram, T., Krey, V., Lamarque, J.-F., et al. (2011). The representative concentration pathways: an overview. *Climatic change*, 109(1-2):5.

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*, 114(38):E7910–E7918.

A Appendix

A.1 Temperature response functions

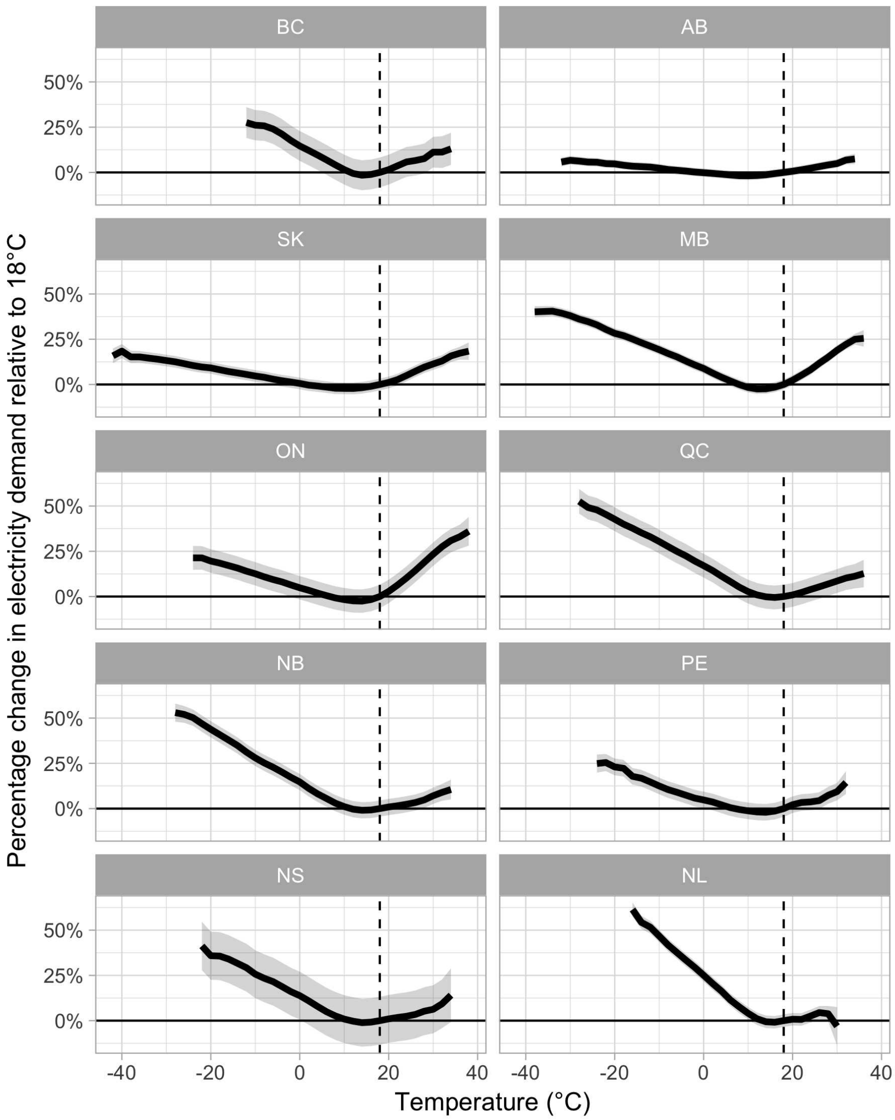

Figure A1: Temperature response functions for each province

Notes: Temperature response functions estimate the percentage change in demand relative to 18°C. Shaded areas represent 95% confidence intervals using Newey-West HAC standard errors.

A.2 Demand projection tables

Table A1: Change in average demand, high emission scenario (RCP8.5)

End-of-century projection

	Without Adaptation (Δ%)			*With Adaptation (Δ%)*		
	Summer	Winter	Annual	Summer	Winter	Annual
BC	1.0	-5.7	-2.1	12.0	-4.8	4.2
AB	1.1	-1.0	0.2	6.3	0.1	3.6
SK	2.3	-2.5	0.1	6.0	-1.7	2.6
MB	2.8	-6.9	-2.0	10.6	-5.4	2.7
ON	6.9	-3.9	2.2	13.6	-3.3	6.3
QC	-0.1	-9.1	-4.6	11.6	-7.8	2.0
NB	-1.3	-8.9	-5.1	10.3	-7.0	1.6
PE	1.5	-4.1	-1.0	5.8	-3.4	1.7
NS	-0.6	-6.6	-3.4	12.2	-4.6	4.3
NL	-4.9	-12.6	-8.9	-0.7	-11.6	-6.4
CAN	**2.4**	**-6.1**	**-1.6**	**11.1**	**-5.0**	**3.6**

Mid-century projection

	Without Adaptation (Δ%)			*With Adaptation (Δ%)*		
	Summer	Winter	Annual	Summer	Winter	Annual
BC	0.1	-3.4	-1.5	6.6	-2.7	2.3
AB	0.6	-0.6	0.1	3.6	0.3	2.2
SK	1.2	-1.4	0.0	2.6	-1.0	1.0
MB	1.2	-3.9	-1.3	3.7	-3.3	0.3
ON	3.8	-2.6	1.0	5.8	-2.3	2.2
QC	-0.4	-5.6	-3.0	4.7	-4.8	0.0
NB	-1.1	-5.7	-3.4	6.2	-4.0	1.1
PE	0.7	-2.7	-0.8	3.4	-2.1	1.0
NS	-0.6	-4.2	-2.3	7.3	-2.5	2.7
NL	-3.2	-8.1	-5.7	-0.8	-7.3	-4.1
CAN	**1.1**	**-3.8**	**-1.2**	**5.0**	**-3.1**	**1.2**

Notes: Summer refers to April–October, winter refers to November–March.

Table A2: Change in average demand, medium emission scenario (RCP4.5)

End-of-century projection

	Without Adaptation (Δ%)			With Adaptation (Δ%)		
	Summer	Winter	Annual	Summer	Winter	Annual
BC	0.1	-3.2	-1.5	7.7	-2.3	3.0
AB	0.5	-0.6	0.1	4.2	0.6	2.6
SK	1.0	-1.3	0.0	3.8	-0.4	1.9
MB	1.1	-3.7	-1.3	6.9	-2.0	2.5
ON	3.5	-2.5	0.9	8.7	-1.8	4.1
QC	-0.5	-5.3	-2.9	8.4	-3.8	2.3
NB	-1.1	-5.3	-3.2	7.4	-3.1	2.1
PE	0.6	-2.5	-0.8	3.8	-1.8	1.3
NS	-0.6	-3.9	-2.2	8.8	-1.7	3.8
NL	-2.9	-7.4	-5.2	0.0	-6.3	-3.3
CAN	**0.9**	**-3.5**	**-1.1**	**7.5**	**-2.3**	**2.9**

Mid-century projection

	Without Adaptation (Δ%)			With Adaptation (Δ%)		
	Summer	Winter	Annual	Summer	Winter	Annual
BC	-0.1	-2.5	-1.2	5.5	-1.7	2.1
AB	0.4	-0.5	0.0	3.0	0.5	1.9
SK	0.8	-1.0	0.0	2.1	-0.5	0.9
MB	0.7	-2.8	-1.0	3.0	-2.1	0.5
ON	2.7	-1.9	0.7	4.5	-1.7	1.8
QC	-0.4	-4.1	-2.2	4.2	-3.2	0.5
NB	-0.9	-4.3	-2.6	5.6	-2.4	1.6
PE	0.5	-2.0	-0.6	2.8	-1.4	0.9
NS	-0.5	-3.1	-1.7	6.5	-1.3	2.8
NL	-2.3	-5.9	-4.2	-0.2	-5.0	-2.7
CAN	**0.7**	**-2.7**	**-0.9**	**4.2**	**-2.0**	**1.3**

Notes: Summer refers to April–October, winter refers to November–March.

Table A3: Change in peak hour demand, high emission scenario (RCP8.5)

End-of-century projection

	Historical		*Without Adaptation*				*With Adaptation*			
	Peak	MW	Peak	MW	ΔMW	Δ%	Peak	MW	ΔMW	Δ%
BC	Winter	11039	Winter	10422	-617	-5.6	Summer	12213	1174	10.6
AB	Winter	11229	Winter	11124	-105	-0.9	Summer	12920	1691	15.1
SK	Summer	4654	Summer	4801	147	3.2	Summer	4942	288	6.2
MB	Winter	4366	Winter	4120	-246	-5.6	Summer	4213	-153	-3.5
ON	Summer	27005	Summer	29817	2812	10.4	Summer	36329	9324	34.5
QC	Winter	39266	Winter	36280	-2986	-7.6	Winter	37390	-1876	-4.8
NB	Winter	3326	Winter	3052	-274	-8.2	Winter	3220	-106	-3.2
PE	Winter	265	Winter	248	-17	-6.4	Summer	250	-15	-5.5
NS	Winter	2192	Winter	2068	-124	-5.7	Summer	2549	357	16.3
NL	Winter	1523	Winter	1393	-130	-8.5	Winter	1416	-107	-7.0

Mid-century projection

	Historical		*Without Adaptation*				*With Adaptation*			
	Peak	MW	Peak	MW	ΔMW	Δ%	Peak	MW	ΔMW	Δ%
BC	Winter	11039	Winter	10698	-341	-3.1	Summer	11046	7	0.1
AB	Winter	11229	Winter	11229	0	0.0	Summer	12090	861	7.7
SK	Summer	4654	Summer	4721	68	1.5	Summer	4753	99	2.1
MB	Winter	4366	Winter	4265	-101	-2.3	Winter	4240	-126	-2.9
ON	Summer	27005	Summer	28757	1752	6.5	Summer	31270	4265	15.8
QC	Winter	39266	Winter	37402	-1864	-4.7	Winter	37964	-1302	-3.3
NB	Winter	3326	Winter	3267	-59	-1.8	Winter	3421	95	2.8
PE	Winter	265	Winter	251	-14	-5.4	Winter	252	-13	-4.9
NS	Winter	2192	Winter	2106	-86	-3.9	Summer	2317	125	5.7
NL	Winter	1523	Winter	1457	-66	-4.3	Winter	1476	-47	-3.1

Notes: Timing of seasonal peak listed in "Peak" columns. Summer refers to Apr–Oct. Winter refers to Nov–Mar.

Table A4: Change in peak hour demand, medium emission scenario (RCP4.5)

End-of-century projection

	Historical		*Without Adaptation*				*With Adaptation*			
	Peak	MW	Peak	MW	ΔMW	Δ%	Peak	MW	ΔMW	Δ%
BC	Winter	11039	Winter	10698	-341	-3.1	Summer	11661	622	5.6
AB	Winter	11229	Winter	11229	0	0.0	Summer	12519	1290	11.5
SK	Summer	4654	Summer	4735	81	1.7	Summer	4801	148	3.2
MB	Winter	4366	Winter	4265	-101	-2.3	Winter	4316	-50	-1.1
ON	Summer	27005	Summer	28637	1632	6.0	Summer	34406	7401	27.4
QC	Winter	39266	Winter	37402	-1864	-4.7	Winter	38605	-661	-1.7
NB	Winter	3326	Winter	3267	-59	-1.8	Winter	3447	121	3.6
PE	Winter	265	Winter	251	-14	-5.4	Winter	253	-12	-4.7
NS	Winter	2192	Winter	2120	-72	-3.3	Summer	2446	254	11.6
NL	Winter	1523	Winter	1457	-66	-4.3	Winter	1483	-40	-2.6

Mid-century projection

	Historical		*Without Adaptation*				*With Adaptation*			
	Peak	MW	Peak	MW	ΔMW	Δ%	Peak	MW	ΔMW	Δ%
BC	Winter	11039	Winter	10816	-223	-2	Winter	11048	9	0.1
AB	Winter	11229	Winter	11229	0	0.0	Summer	11967	738	6.6
SK	Summer	4654	Summer	4704	51	1.1	Summer	4724	70	1.5
MB	Winter	4366	Winter	4305	-61	-1.4	Winter	4305	-61	-1.4
ON	Summer	27005	Summer	28344	1339	5.0	Summer	30721	3716	13.8
QC	Winter	39266	Winter	38195	-1071	-2.7	Winter	38780	-486	-1.2
NB	Winter	3326	Winter	3267	-59	-1.8	Winter	3425	99	3.0
PE	Winter	265	Winter	251	-14	-5.4	Winter	252	-13	-4.9
NS	Winter	2192	Winter	2139	-53	-2.4	Summer	2281	89	4.1
NL	Winter	1523	Winter	1458	-65	-4.3	Winter	1482	-41	-2.7

Notes: Timing of seasonal peak listed in “Peak” columns. Summer refers to Apr–Oct. Winter refers to Nov–Mar.

A.3 Testing goodness-of-fit of selected observables

Table A5 summarises the right-side and left-side slopes for each province's short run temperature response functions by re-estimating them using only cooling and heating degrees (with a baseline of 14°C) rather than temperature bins. These are presented alongside mean values for the key observables.

Table A5: Temperature response function slope coefficients and observable averages

	RHS slope (CD)	LHS slope (HD)	AC	Electric Heat	Res Share
AB	0.005	0.002	0.195	0.399	0.165
BC	0.007	0.013	0.205	0.477	0.317
MB	0.013	0.010	0.632	0.548	0.355
NB	0.006	0.014	0.291	0.707	0.398
NL	0.006	0.021	0.060	0.639	0.371
NS	0.005	0.012	0.178	0.451	0.382
ON	0.016	0.007	0.753	0.417	0.343
PE	0.008	0.007	0.249	0.321	0.142
QC	0.006	0.013	0.509	0.820	0.354
SK	0.008	0.004	0.560	0.218	0.172

Notes: The RHS slope and LHS slope refer to the right- and left-side slope coefficients for the short run temperature response functions, i.e. the change in log(demand) for a 1°C change in temperature when above and below 14°C, respectively. AC and Electric Heat are the mean penetration of air conditioners and electric heating systems per household by province over the 2001-2015 period. Res Share represents the share of total demand attributed to the residential sector.

Table A6 presents the results of regressions to determine the explanatory power of the observables in explaining the provincial heterogeneity in slopes of short run temperature response functions. Specifically, we regress, separately, the left- and right-side slopes of the short run temperature response functions against provincial mean values of the observables, along with variants that interact residential shares with air conditioner and electric heat penetration. The explanatory power of these three observables is strong, showing an R-squared of 77-85% depending on specification. Despite only 10 observations (one for each province) the adjusted R-squared remains relatively strong despite only three observable variables used to explain the provincial heterogeneity. Figure A2 presents added-variable residual plots (for specification (i) in Table A6) to demonstrate the relationship between the RHS slope of the temperature response functions and the three observable variables.

Table A6: Testing for the explanatory power of observables

	Dependent variable:					
	Cooling Degree Slope			Heating Degree Slope		
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
AC	0.012***	−0.014	−0.014	−0.010*	0.017	0.022
	(0.003)	(0.015)	(0.030)	(0.005)	(0.024)	(0.047)
Electric Heat	−0.011	−0.015*	−0.016	0.006	0.011	0.021
	(0.006)	(0.006)	(0.050)	(0.009)	(0.010)	(0.078)
AC:Res Share		0.078	0.080		−0.081	−0.095
		(0.045)	(0.089)		(0.070)	(0.139)
Res Share:Electric Heat			0.003			−0.026
			(0.133)			(0.207)
Res Share	0.017	−0.002	−0.004	0.034*	0.054*	0.067
	(0.011)	(0.015)	(0.071)	(0.016)	(0.023)	(0.110)
Observations	10	10	10	10	10	10
R^2	0.765	0.854	0.854	0.781	0.827	0.828
Adjusted R^2	0.648	0.737	0.671	0.671	0.689	0.613
df	6	5	4	6	5	4

Notes: Standard errors clustered by year-month. *p<0.1; **p<0.05; ***p<0.01

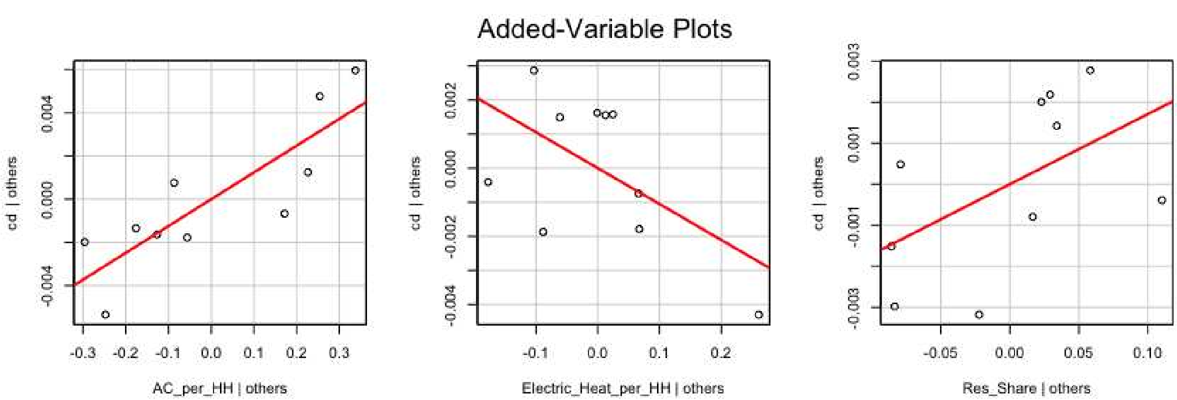

Figure A2: Added-value plot for Model 1 (above)

A.4 Robustness checks for air conditioner regression

Table A7 shows the results from estimating a linear probability version of Eq.8. The table reports the results from four separate regressions, each using an alternative definition of the hot temperature variable $\tilde{T}$ in each city, as described above. In each case, the table shows that there is a positive and highly statistically significant relationship between the prevalence of warm weather and the penetration of air conditioners. We focus on column (4), because it provides the best fit and is a natural way to describe the relationship between air conditioner penetration and the climate. In this column, we regress air conditioner penetration on the average July and August daily maximum temperature, as well as other covariates as described above. The table shows that for each 1 degree Celsius increase in the maximum daily July-August temperature, the penetration of residential air conditioners increases by 16.8 percentage points. Air conditioner penetration is clearly quite sensitive to the prevalence of warm summer temperatures.

Table A7: Linear probability model for air conditioner penetration with alternative climatic variables

	Dependent variable:			
	ac			
	(1)	(2)	(3)	(4)
highestMonthlyMean	0.145***			
	(0.001)			
highestTemp		0.111***		
		(0.001)		
meanJulyAug			0.142***	
			(0.001)	
meanmaxJulyAug				0.168***
				(0.001)
Household weights	No	No	No	No
Observations	33,591	33,591	33,591	33,591
R^2	0.304	0.280	0.253	0.325
Adjusted R^2	0.304	0.279	0.253	0.325
Residual Std. Error (df = 33575)	0.413	0.421	0.428	0.407
F Statistic (df = 15; 33575)	978.756***	869.343***	758.589***	1,077.589***
Note:				*p<0.1; **p<0.05; ***p<0.01

In Table A8, we provide some robustness checks for this main result. In column (1), we estimate the same model, but this time using sampling weights provided by Statistics Canada to ensure the sample is representative. Not surprisingly, the results are not substantially affected. In column (2), we estimate a probit model rather than a linear probability model. The key coefficient remains positive and highly statistically significant, and the average marginal effect remains very close to the estimate using the linear probability model: a one degree increase in the mean daily maximum July-August temperature is projected to increase air conditioner penetration by 14.9 percentage points. In column (3), we estimate a linear probability model with province fixed effects. In this case, the effect of climate on air conditioner penetration is identified from within-province variation in temperature, which eliminates any province-specific unobserved variables, such as building regulations or norms. The effect of climate is somewhat smaller in this specification, but remains significant and relatively close to the original specification.

Table A8: Alternative functional forms for air conditioner penetration

	Dependent variable:		
	ac		
	OLS	*probit*	*OLS*
	(1)	(2)	(3)
meanmaxJulyAug	0.162***	0.527***	0.115***
	(0.001)	(0.006)	(0.002)
Average marginal effect	-	0.149 (0.001)	-
Household weights	Yes	No	No
Province FEs	No	No	Yes
Observations	33,591	33,591	33,591
R^2	0.306		0.365
Adjusted R^2	0.306		0.365
Log Likelihood		−16,811.600	
Akaike Inf. Crit.		33,655.200	
Residual Std. Error	12.434 (df = 33575)		0.395 (df = 33567)
F Statistic	986.617*** (df = 15; 33575)		839.626*** (df = 23; 33567)
Note:			*p<0.1; **p<0.05; ***p<0.01