

Climate Amenities, Climate Change, and American Quality of Life

David Albouy, Walter Graf, Ryan Kellogg, and Hendrik Wolff*

April, 2015

Abstract

We present a hedonic framework to estimate U.S. households' preferences over local climates, using detailed weather and 2000 Census data. We find that Americans favor a daily average temperature of 65 degrees Fahrenheit, that they will pay more on the margin to avoid excess heat than cold, and that damages increase less than linearly over extreme temperatures. These preferences vary by location due to sorting or adaptation. Changes in climate amenities under business-as-usual predictions imply annual welfare losses of 1 to 4 percent of income by 2100, holding technology and preferences constant. (JEL H49, I39, Q54, R10)

* Albouy: Department of Economics, University of Illinois, and NBER (email: albouy@umich.edu); Graf: Department of Agricultural & Resource Economics, University of California, Berkeley (email: walter.graf@berkeley.edu); Kellogg: Department of Economics, University of Michigan, and NBER (email: kellogr@umich.edu); Hendrik Wolff, Department of Economics, University of Washington (email: hgwolff@uw.edu). Acknowledgments: Albouy acknowledges financial support from NSF grant SES-0922340. We thank Kenneth Chay, Don Fullerton, Philip Haile, Kai-Uwe Kühn, Matthew Turner, Wolfram Schlenker, the editor Daniel Phaneuf, and three anonymous referees for particularly helpful criticisms and suggestions. We are also grateful for comments from seminar participants at Boston, Brown, Calgary, Columbia, CU-Boulder, Duke, EPA, Harvard, Houston, Illinois, Maryland, Michigan, Michigan State, Minnesota, NYU, Olin, Oregon, Princeton, RFF, Rice, Stanford, Texas A&M, TREE, UC Berkeley, UCLA, UNR, Washington, Western Michigan, and Wisconsin, as well as conference attendees at the AEA, Cowles, NASM, NARSC, NBER EEE, RSQE, and WCERE.

The chemistry of the human body makes our health and comfort sensitive to climate. Every day, climate influences human activity, including diet, chores, recreation, and conversation. Households spend considerable amounts on housing, energy, clothing, and travel to protect themselves from extreme climates and to enjoy comfortable moderation. Geographically, climate impacts the desirability of different locations and the quality of life they offer; few seek to live in the freezing tundra or oppressively hot deserts. Given the undeniable influence climate has on economic decisions and welfare, we seek to estimate the dollar value American households place on climate amenities, including sunshine, precipitation, humidity, and especially temperature.

Valuing climate amenities not only helps us to understand how climate affects welfare and where people live, but also helps to inform policy responses to climate change. Global climate change threatens to alter local climates, most obviously by raising temperatures. *A priori*, the welfare impacts of higher temperatures are ambiguous: households may suffer from hotter summers but benefit from milder winters. Ultimately, these impacts depend on where households are located, the changes in climate amenities they experience, and how much they value these changes.

In this paper, we estimate the value of climate amenities in the United States by examining how households' willingness to pay (WTP) to live in different areas varies with climate in the cross-section. Following the intuition laid out by Rosen (1974, 1979) and Roback (1982), and later refined by Albouy (2011), we measure WTP by developing a local quality-of-life (QOL) index based on how much households pay in costs-of-living relative to the incomes they receive. The U.S. is a particularly appropriate setting in which to use this method as it has a large population that is mobile over areas with diverse climates. Globally, the U.S. lies in a temperate zone, with some areas that are quite hot (Arizona) while others are quite cold (Minnesota), and some with extreme seasonality (Missouri) while others are mild year-round (coastal California). This variation allows us to identify preferences over a broad range of habitable climates.

We adopt this hedonic approach as there are no explicit markets for climate amenities, only an implicit market based on household location choices. Our estimates of amenity

values primarily reflect impacts of exposure to climate on comfort, activity, and health, including time use (Graff Zivin and Neidell 2012) and mortality risk (Barreca *et al.* 2013, Deschênes and Greenstone 2011). They exclude costs from residential heating, cooling, and insulation. As such, the value of the climate amenities we estimate does not appear in national income accounts, and so neither would the impact of climate change on these amenities. Our study therefore complements work that assesses how climate directly affects national income through agricultural and urban productivity (for a survey of the climate and productivity literature, see Tol 2002 and 2009).

We adopt a cross-sectional estimation strategy in the tradition of Mendelsohn *et al.* (1994) rather than a time-series panel approach for several reasons. First, yearly changes in weather are unlikely to affect households' WTP to live in an area: WTP should depend on long-run climate rather than the outcome of the weather in the most recent year. Second, low frequency changes are not very informative as long-run secular climate changes so far have been slight, particularly relative to changes in technology—especially air conditioning (Barreca *et al.* 2013)—and local economic conditions. Third, Kuminoff and Pope (2012) have shown that temporal changes in the capitalization of amenities do not typically translate to measures of WTP. Finally, households can mitigate potential damages from climate through adaptation—say, by insulating homes, changing wardrobes, or adopting new activities—which cross-sectional methods account for, thereby making our estimates more relevant for assessing the impact of climate change.

An unavoidable drawback of our estimation strategy is that it requires climate amenities to be uncorrelated with the influence unobserved local amenities have on QOL. This untestable assumption is hard to circumvent, as there do not appear to be any viable instrumental variables for climate. Instead, we examine the potential for omitted viable bias by testing the robustness of our estimates to an array of specifications and powerful controls, following the intuition laid out by Altonji *et al.* (2005) and the literature on agricultural yields and farmland values (Schlenker *et al.* 2006, Deschênes and Greenstone 2007, Schlenker and Roberts 2009, Deschênes and Greenstone 2012, and Fisher *et al.* 2012). While the stability of our estimates across these specifications is reassuring, we

nonetheless acknowledge that our research design cannot conclusively rule out the possibility that unobserved factors are influencing our estimates.

Prior hedonic studies investigating U.S. climate preferences have been few and conflicting. Estimates of WTP for incremental warming range from positive (Hoch and Drake 1974, Moore 1998), to approximately zero (Nordhaus 1996), to negative (Cragg and Kahn 1997 and 1999, Kahn 2009, Sinha and Cropper 2013).¹ Our paper makes three contributions to the literature on climate amenities by drawing from recent innovations in the literatures on climate damages to agriculture and health, quality of life measurement, and hedonic estimation under preference heterogeneity. First, we characterize climates using the full distribution of realized daily temperatures rather than seasonal or monthly averages, allowing us to explore how households value progressively extreme temperatures. Prior research into climate amenities has ignored the importance of flexibly modeling temperature profiles, even though other research has shown that extreme temperatures, not averages, are especially harmful to crop yields (Schlenker *et al.* 2006, Schlenker and Roberts 2009) and health (Barreca 2012, Deschênes and Greenstone 2011). Second, following Albouy and Lue (forthcoming), our QOL estimates account for commuting costs, local expenditures beyond housing, and federal taxes on wages. We show that previous work, by ignoring these factors, produced noisy and misleading estimates of climate valuations, while our results are robust to many alternative specifications, such as including controls for every state. Third, we apply methods by Bajari and Benkard (2005) to model unobserved heterogeneity in households' preferences, thereby allowing for sorting based on differences in (dis)taste for cold or heat and for adaptation to local climates.

Our estimates of climate preferences yield four main results. First, we find that Americans most prefer daily average temperatures—the average of the daily high and daily

¹ Hedonic studies focusing on countries other than the U.S. include Maddison and Bigano (2003) in Italy, Rehdanz (2006) in Great Britain, Rehdanz and Maddison (2009) in Germany, and Timmins (2007) in Brazil. In addition, alternative non-hedonic approaches have been used to estimate the impact of climate change on amenities. Shapiro and Smith (1981) and Maddison (2003) use a household production function approach, Fritjers and Van Praag (1998) use hypothetical equivalence scales, and Rehdanz and Maddison (2005) link a panel of self-reported happiness across 67 countries to temperature and precipitation.

low—near 65 degrees Fahrenheit (°F), agreeing with standard degree-day models that predict little need for heating or cooling at this temperature. Second, on the margin, households pay more to avoid a degree of excess heat than a degree of excess cold. Third, we find that the marginal WTP (MWTP) to avoid extreme temperatures is not substantially greater than the MWTP to avoid moderate heat and cold. Put another way, households will pay more to turn a moderately cold day into a perfect day than to turn a bitterly cold day into a moderately cold day. This finding is consistent with evidence that households reduce their time outdoors as temperatures become uncomfortable, reducing their sensitivity to further temperature changes (Graff Zivin and Neidell 2012). Fourth, we find that households in the South are particularly averse to cold. This result is consistent with models of both residential sorting and adaptation. Conversely, we do not find evidence that southern households are less heat-averse than northern households.

We apply our estimated climate preferences by simulating how climate change may impact welfare by improving or reducing the value of climate amenities. For our climate change predictions, we rely primarily on the business-as-usual A2 scenario used in the Intergovernmental Panel on Climate Change (IPCC) fourth assessment report (2007), which predicts a 7.3°F increase in U.S. temperatures by 2100. Our simulated welfare effects are predicated on technology and preferences remaining constant and are therefore best interpreted as a benchmark case. This assumption is common to most estimates of climate change damages, including the agricultural and health literatures referenced above. We view endogenizing technology and preferences as beyond the scope of this paper's climate change application, and we leave this issue for future work.

We project that on average, Americans would pay 1 to 4 percent of their annual income to avoid predicted climate change. While damages are rather severe in the South, we find that most areas in the North also suffer because: (1) they lose many pleasant summer days in exchange for only moderately warmer winter days, and (2) Northerners are willing to pay less to reduce cold than are Southerners. Welfare impacts are reduced slightly when we model migration responses.

Below, we first discuss the model we use to derive local QOL estimates from wages and housing costs. Section II describes the current and projected climate data, and section III presents estimates of climate preferences assuming preferences are homogenous. These results provide a baseline for section IV, which allows for heterogeneous preferences. Section V then discusses the implications of our preference estimates for climate change impacts. Section VI considers alternative models of QOL, and section VII concludes.

I. Hedonic Estimates of Quality of Life

The intuition underlying our approach is that households pay higher prices and accept lower wages to live in areas with desirable climate amenities. Below we discuss the hedonic framework underlying this intuition, how we calculate wage and cost of living differences across areas, and how we combine these to create a single-index QOL measure for each location. Our approach is rooted in the conceptual framework of Rosen (1974, 1979) and Roback (1982) and adopts important recent contributions to this framework from the hedonic literature. In particular, we follow Albouy (2011) and Albouy and Lue (forthcoming) to properly weight wages and housing prices when creating the QOL measure, and we adopt Bajari and Benkard (2005) to allow for preference heterogeneity. These innovations ultimately prove to be consequential in obtaining preferences for climate, as we discuss further below.

A. A model of QOL using local cost of living and wage differentials

The U.S. economy consists of locations, indexed by j , which trade with each other and share a population of perfectly mobile, price-taking households, indexed by i (we discuss estimates that relax the perfect mobility assumption in section VI). These households have preferences over two consumption goods: a traded numeraire good x and a non-traded "home" good y , with local price p_j that determines local cost of living. Households earn wage income w_j^i that is location-dependent and own portfolios of land and capital

that pay a combined income of R^i . Gross household income is $m_j^i = R^i + w_j^i$, out of which households pay federal taxes of $\tau(m_j^i)$. Federal revenues are rebated lump sum.²

Each location is characterized by a K -dimensional vector of observable amenities \mathbf{Z}_j , including climate, and a scalar characteristic ξ_j that is observable to households but not econometricians.³ We assume a continuum of locations so that the set of available characteristics (\mathbf{Z}, ξ) is a complete, compact subset of \mathbb{R}^{K+1} .⁴ Each household i seeks out the location j that maximizes its utility, given by $u_j^i = V^i(p_j, w_j^i, \mathbf{Z}_j, \xi_j)$. This indirect utility function is assumed to be continuous and differentiable in all its arguments, strictly increasing in w^i and ξ_j , and strictly decreasing in p_j . Households are permitted to have heterogeneous preferences over (\mathbf{Z}, ξ) .

On the supply side, we assume that firms face perfectly competitive input and output prices and earn zero profits, offering higher wage levels in more productive locations. We model each household i 's wage in location j as $\varphi^i w_j$, where φ^i is household-specific skill and w_j is the local wage level.⁵

Let $\mathbf{p}(\mathbf{Z}_j, \xi_j)$ and $\mathbf{w}(\mathbf{Z}_j, \xi_j)$ denote the functions relating wages and prices to local characteristics. These functions are determined in equilibrium by households' demands for local amenities, firms' location decisions, and the availability of land.⁶ On the demand

² We also apply adjustments for state taxes and tax-benefits to owner occupied housing, discussed in Albouy (2011), which prove to be minor in practice.

³ This set up omits an idiosyncratic unobserved preference shock ε_{ij} from households' utility function. Relaxing this assumption implies that the QOL measure may depend not just on prices and wages but also on population levels or changes. We allow for such dependence in some of our specifications as discussed below (for example, by adjusting our wage estimates for migration per Dahl (2002)), none of which substantially alters our conclusions. See Albouy (2011) for a more general discussion of addressing idiosyncratic preferences in QOL estimates.

⁴ In our empirical implementation, this completeness assumption is an approximation, since we examine a finite, though large, number of U.S. PUMAs. We discuss how this assumption can be relaxed in section IV and Appendix 4 via a set identification procedure.

⁵ By using a single index of skill, we abstract away from the possibility that some households have a comparative advantage in certain locations. Relaxing this assumption has implications similar to those for allowing an idiosyncratic unobserved preference shock (see footnote 8 above).

⁶ The existence and differentiability of $\mathbf{p}(\mathbf{Z}_j, \xi_j)$ and $\mathbf{w}(\mathbf{Z}_j, \xi_j)$ hold under mild conditions. Rosen demonstrates the existence and differentiability of the quality of life function $\hat{Q}(\mathbf{Z}_j, \xi_j)$ under the assumptions on demand given here and perfectly competitive land supply, while Bajari and Benkard (2005) show that a Lipschitz condition on $V^i(p_j, w_j^i, \mathbf{Z}_j, \xi_j)$ is sufficient even under imperfect competition. Given this result, the separate existence and differentiability of $\mathbf{p}(\mathbf{Z}_j, \xi_j)$ and $\mathbf{w}(\mathbf{Z}_j, \xi_j)$ are given by the separate

side, household utility maximization implies the following first-order condition for each characteristic k :

$$(1) \quad \frac{1}{m_j^i \lambda^i} \frac{\partial V^i(p_j, w_j, \mathbf{Z}_j, \xi_j)}{\partial Z_k} = \frac{p_j y_j^i}{m_j^i} \frac{\partial \ln \mathbf{p}(\mathbf{Z}_j, \xi_j)}{\partial Z_k} - \frac{w_j^i (1 - \tau')}{m_j^i} \frac{\partial \ln \mathbf{w}(\mathbf{Z}_j, \xi_j)}{\partial Z_k}$$

where λ^i is the marginal utility of income and τ' is the average marginal tax rate on labor income. This equation relates household i 's marginal valuation of characteristic k , as a fraction of income, to differential changes in the logarithms of the cost of living and wage differentials at j .

Operationally, we develop a QOL index to indicate the willingness-to-pay of households, averaged by income, from the right-hand side of (1). This measure at j , denoted \hat{Q}_j , is a weighted combination of \hat{p}_j and \hat{w}_j , the differentials in log housing-costs and wages relative to the U.S. income-weighted average, according to the formula

$$(2) \quad \hat{Q}_j = s_y \hat{p}_j - (1 - \tau') s_w \hat{w}_j = 0.33 \hat{p}_j - 0.50 \hat{w}_j$$

where s_y denotes the average share of income spent on local goods and s_w the average share of income from wages. The second equality substitutes in values for these parameters of $s_y = 0.33$, $s_w = 0.75$, and $\tau' = 0.33$. For additional details, including the incorporation of local non-housing expenditures into s_y , see Albouy (2011). Section VI considers estimates allowing for heterogeneity in s_y and s_w .

Let $\hat{Q}(\mathbf{Z}_j, \xi_j)$ denote QOL as a function of local characteristics, per (2) and the functions $\mathbf{p}(\mathbf{Z}_j, \xi_j)$ and $\mathbf{w}(\mathbf{Z}_j, \xi_j)$. Then, by condition (1), for any household i in j , the marginal willingness to pay (MWTP) for characteristic k is equal to the derivative of $\hat{Q}(\mathbf{Z}_j, \xi_j)$ with respect to k :

$$(3) \quad \frac{1}{m_j^i \lambda^i} \frac{\partial V^i(p_j, w_j, \mathbf{Z}_j, \xi_j)}{\partial Z_k} = \frac{\partial \hat{Q}(\mathbf{Z}_j, \xi_j)}{\partial Z_k}$$

mobility conditions on households and firms and an assumption that local productivity differentials are continuous and differentiable in \mathbf{Z}_j and ξ_j .

Condition (3) is illustrated in Figure 1 in the case of a single characteristic, average summer temperature T_s . The bold line denotes a hypothetical function $\widehat{Q}(T_s)$ that is decreasing in T_s , indicating that milder weather is “paid for” through either higher housing prices or lower wages. The slope of $\widehat{Q}(T_s)$ at any given location is the hedonic price for a marginal increase in temperature at that location. This hedonic price is equal to households’ MWTP for T_s at that location, as shown for locations A and B on the figure. As shown, households at A have a higher MWTP to avoid heat than do households at B, consistent with sorting or adaptation.

B. Estimates of wage and housing cost differentials, and QOL

We follow Albouy (2011) to estimate wage and housing-cost differentials using the 5 percent sample of Census data from the 2000 Integrated Public Use Microdata Series (IPUMS). Geographically, these data are available by Public Use Microdata Areas (PUMAs), which contain populations of 100,000–300,000 and form the main unit of our analysis. We summarize this procedure below; for more details see Albouy (2011).

We calculate wage differentials by PUMA, \widehat{w}_j , using the logarithm of hourly wages for full-time workers aged 25 to 55 and controlling for observable skill and occupation differences across workers. Specifically, we regress the log wage of worker i on PUMA-indicators μ_j^w and extensive controls X_i^w (each interacted with gender) for education, experience, race, occupation, and industry, as well as veteran, marital, and immigrant status, in an equation of the form $\ln w_{ij} = X_i^w \beta^w + \mu_j^w + \varepsilon_{ij}^w$. The estimates of the μ_j^w are used as the PUMA wage differentials \widehat{w}_j following a refinement, per Albouy and Lue (forthcoming), so that they reflect wages by place-of-work rather than place-of-residence, netting out differences in commuting costs.

Our model interprets the \widehat{w}_j as the causal effect of a PUMA’s characteristics on a worker’s wage, while the observable and unobservable skill differences across workers, the X_i^w and ε_{ij}^w , are an analogue to the φ^i factors in the model. This interpretation requires that any sorting of workers across locations based on unobserved skills or a spatial

match component of wages does not substantially affect observed wage premia. This assumption receives mixed support in the literature. Glaeser and Maré (2001) and Baum-Snow and Pavan (2012) find that unobserved skill and match-based sorting contribute little to city-size wage premia; however, Gyourko, Mayer, and Sinai (2013) find that a select group of “superstar cities” may disproportionately attract high-skilled workers. Dahl (2002) finds that selective migration biases estimates of the returns to education (though not the range of returns across states), and Kennan and Walker (2011) finds a role for location-specific job matches in migration decisions.

In light of the uncertainty in the literature, we address the potential for skill and match-based sorting in a series of alternative specifications. First, we adopt the method used in Dahl (2002) to adjust our wage estimates \hat{w}_j for selective migration by including a flexible control function of migration probabilities in our wage equation. Second, in a closely related specification, we directly adjust our QOL estimates for each PUMA using the PUMA’s rate of net in-migration between 1990 and 2000 (in percent) and a mobility cost estimate from Notowidigdo (2012). Finally, we guard against effects from “superstar cities” by estimating a specification in which superstar metropolitan areas (as defined in Gyourko, Mayer, and Sinai (2013)) are dropped from the sample. These specifications are described in more detail in section VI below and ultimately yield estimates of climate preferences and welfare effects that are not qualitatively different from our baseline estimates.

To calculate housing cost differentials, we use housing values and gross rents, including utilities. Following previous studies, we convert housing values to imputed rents at a discount rate of 7.85 percent (Peiser and Smith 1985) and add in utility costs to make them comparable to gross rents. This approach follows the standard practice in the QOL literature from Blomquist *et al.* (1988) to Chen and Rosenthal (2008) and is required by the data because utility costs are included in gross rents. We then calculate housing-cost differentials with a regression of rents on flexible controls Y_{ij} (each interacted with renter status) for size, rooms, acreage, commercial use, kitchen and plumbing facilities, type and age of building, and the number of residents per room. This regression takes the form

$\ln p_{ij} = Y_{ij}\beta^p + \mu_j^p + \varepsilon_{ij}^p$. The estimates of the μ_j^p are then used as PUMA-level housing cost differentials \hat{p}_j . Proper identification of housing-cost differences requires that they do not vary systematically with unobserved housing quality across locations.

We incorporate energy and insulation costs in our housing-cost measure because doing so allows us to interpret our QOL differentials as solely reflecting the value of nonmarket climate amenities rather than the effect of climate on utility costs. Hence, our QOL differentials will reflect the disamenity of outdoor exposure to climate and the disamenity of adverse indoor temperatures to the extent they are not completely mitigated by insulation and energy use. In addition, the QOL estimates will incorporate any disamenity from spending more time indoors to avoid uncomfortable outdoor temperatures.

Descriptive statistics for QOL are given at the bottom of Table 1,⁷ and QOL differentials across PUMAs for the year 2000 are mapped in Figure 2. These estimates show that households find the amenities in urban areas, coastal locations, and certain mountain areas to be quite desirable. Areas in the middle of the country, where seasons are more extreme, tend to be less desirable, although the variation is considerable. As discussed in Albouy (2011), our QOL estimates correlate well with non-economic measures of QOL, such as the “livability” rankings in the *Places Rated Almanac* (Savageau 1999). Moreover, the QOL model correctly predicts the relationship between housing costs and wages, controlling for observable amenities.

II. Data

We estimate our main specifications at the PUMA-level using 2,057 PUMAs covering the contiguous 48 states as of the 2000 census.⁸ In this section, we summarize our dataset,

⁷ The mean QOL differential is not exactly zero in Table 1 because the table shows unweighted data, while QOL differentials are defined so that the income-weighted mean is zero.

⁸ We have also aggregated our data to the MSA-level and run some of our empirical specifications at an MSA-level resolution. The point estimates for preferences and climate change welfare impacts are similar to those discussed below, with modestly larger standard errors (see Appendix Table A1.1, column R13). We believe that the MSA-level results are less precise because MSAs are frequently too large to capture important micro-climates, particularly in densely-populated coastal areas such as San Francisco.

covering recent historical climate, climate-change projections, and other variables. Additional details are provided in Appendix 2.

A. Recent historical climate data

Our daily average temperature data over 1970–1999 originate from Schlenker and Roberts (2009), which provides historic daily temperatures at a 4 kilometer-square resolution. From these data, we create temperature bins with a width of 0.9°F (0.5°C) and calculate the average number of days at each 4 kilometer grid point, over 1970–1999, in which the daily average temperature—calculated as the mean of the daily high and low—fell within each bin.⁹ Within each PUMA, we average the bin distribution at each grid point to yield a PUMA-level dataset.

We obtain monthly precipitation and humidity data from the Parameter elevation Regressions on Independent Slopes Model (PRISM) at the same grid points for 1970–1999, averaging them by month-of-year at the PUMA-level. We obtain sunshine data, measured as the percent of daylight hours for which the sun is not obscured by clouds, by month-of-year from 156 weather stations from the National Climactic Data Center. We interpolate PUMA-level data on sunshine from the four closest weather stations.¹⁰

⁹ As an alternative specification, we fit a sinusoidal curve to each day’s high and low temperature and then find the fraction of each day that falls within each temperature bin, following Schlenker and Roberts (2009). The resulting U.S. population-weighted average temperature distribution is plotted in appendix Figure A1.3 and is necessarily broader than that obtained from daily means. Figure A1.3 also presents estimated WTPs for temperature that use the procedure discussed in section III and are analogous to Figure 4 (7th degree spline) and Figure 5 (panel B) in the main text. The estimated WTPs using intra-day temperature variation are similar to those in the main text: WTP to avoid heat is greater than WTP to avoid cold, and the WTP curves flatten at extreme temperatures. This flattening occurs at farther extremes than in the main specification, since extreme low and high temperatures are necessarily lower and higher, respectively, than extreme daily mean temperatures. Figure A1.3 does not present an analogous distribution of future climate because our climate model data source does not provide high and low temperature information.

¹⁰ Our sunshine data are therefore not as fine-scaled as the other variables in our regressions and may incorporate prediction error. This concern is at least partially mitigated by the fact that variation in sunshine is smooth relative to the density of sunshine-measuring weather stations, as shown in appendix Figures A2.1 and A2.2, though microclimates may nonetheless not be captured. Fine-scale variation in precipitation and humidity may help to proxy for finer changes in local sunshine, as they are strongly correlated at coarser levels of geography.

B. Projected climate data

Predicted climate change data (temperatures, precipitation, humidity, and sunshine) are from the third release of the Community Climate System Model. These data were also used in the Intergovernmental Panel on Climate Change Assessment Report 4 (IPCC AR4) released in 2007. We use two business-as-usual scenarios in which no actions to reduce greenhouse gas emissions are taken: the A2 scenario and the A1FI scenario. In both models, data are provided at a resolution of 1.4 degrees longitude (120 km) by 1.4 degrees latitude (155 km), so we interpolate these data to the PUMA level.¹¹ The A2 scenario predicts average (population-weighted) U.S. warming of 7.3°F from the baseline (1970-1999) to the end of the century (2070-2099), while the A1FI scenario predicts warming of 8.6°F. We focus on the A2 scenario but also provide results for A1FI.

Data for present and projected (A2 model) climate variables are given in the top half of Table 1, which summarizes temperature distributions using annual heating degree days (HDD) and cooling degree days (CDD) statistics. HDD is a measure of how cold a location is: it equals the sum, over all days of the year in which it is colder than 65°F, of the difference between 65°F and each day's mean temperature. CDD, a measure of heat, is defined similarly for temperatures greater than 65°F. HDD and CDD are often used by engineers as predictors of heating and cooling loads, and since Graves (1979) they have often been used as measures of climate amenities.

According to the CCSM model predictions, climate change will be manifest primarily through changes in temperature, as seen in Table 1. The average U.S. PUMA will see its number of HDDs fall by 32% and its number of CDDs rise by 97% under the A2 scenario. In contrast, changes to precipitation, relative humidity, and sunshine are predicted to be minor on average.¹² The predicted temperature changes vary considerably by geography, as shown in Figure 3. While substantial increases in both January and July temperatures are predicted nationwide, the interior South is predicted to experience a particularly large increase in days for which the average temperature exceeds 90°F.

¹¹ This interpolation may fail to capture localized changes, particularly in coastal areas.

¹² Some areas are predicted to have appreciable changes, though these average out to be small. Also note that our study is meant to capture the impact of precipitation on QOL, not water supply.

C. Other variables

Table 1 also presents data on the control variables in our econometric specifications. The geographic controls, used in all estimates, include the minimum distances from each PUMA’s centroid to an ocean and Great Lakes coastline, as well as the average slope of the land, to measure hilliness. Demographic data include measures of population density,¹³ educational attainment, age, and racial-ethnic composition. Table A1.2 in the appendix provides evidence on correlations of these geographic and demographic control variables with HDD and CDD. Table 1 also provides statistics on the distribution of population across PUMAs and on the QOL measure.

III. Estimation of WTP for Climate under Homogenous Preferences

A. Specification

We begin by estimating a simple hedonic model in which we assume that climate preferences are homogenous across the U.S. population and that factors (including climate) affecting QOL enter linearly. While this model is highly restrictive, it provides an intuitive introduction to our approach that resembles the previous literature, and it provides a benchmark against which we can later compare a model that allows for preference heterogeneity.

We estimate the impact of marginal changes to climate on QOL using an OLS regression of each PUMA j ’s QOL differential \hat{Q}_j on vectors of climate variables \mathbf{X}_j and other local characteristics \mathbf{D}_j :

$$(4) \quad \hat{Q}_j = \beta \mathbf{X}_j + \gamma \mathbf{D}_j + \xi_j$$

¹³ We use population density rather than population because PUMAs are drawn to have similar populations. Our population density measure is “weighted” in the sense that, within each PUMA, we calculate the population density of each of its census tracts and then take a population-weighted average of these densities. This weighted density gives a better sense of the population density immediately around individuals than a conventional unweighted measure.

The parameters β and γ represent the WTP of households for an additional unit of each element of \mathbf{X}_j and \mathbf{D}_j respectively, measured as a fraction of income. The disturbance term ξ_j is a vertical location characteristic that is observed by households but not by the econometrician. We face two substantial challenges in estimating (4). First, we must select a functional form for how the climate variables – and in particular the temperature distribution – enter into \mathbf{X}_j . Our goal is to use a form that is both flexible and capable of providing precise estimates.¹⁴ Second, consistent estimation of β and γ requires that unobserved factors ξ_j be uncorrelated with \mathbf{X}_j and \mathbf{D}_j . In the absence of instrumental variables, orthogonality between ξ_j and \mathbf{X}_j is a necessary assumption that cannot be tested. We therefore assess the reliability of our estimates of β by studying their sensitivity to a variety of alternative specifications. Ultimately, we view the estimates’ robustness across these regressions as arguing in favor of their interpretation as preference parameters, but it is important to caveat that these results cannot definitively rule out the presence of confounding factors.

We model the WTP for exposure to an additional day at temperature t , relative to a day at 65°F, as the function $f(t)$.¹⁵ 65°F is a natural normalization point because conventional HDD and CDD calculations treat 65°F as a “bliss point” at which neither heating nor cooling is required. If discomfort from heat and cold, in terms of WTP, follow HDD and

¹⁴ We could in principle model rainfall, humidity, and sunshine as flexible distributions, as we do with temperature. We instead simply include these variables as linear regressors, for two reasons. First, these three variables are highly collinear even when entered linearly; allowing for substantial nonlinearities only exacerbates this problem. Second, we wish to focus primarily on temperatures, given that in our climate change application only temperatures change substantially. We have estimated a binned specification for precipitation (6 bins) and humidity (10 bins), similar to that in Barreca (2012); this specification yields similar results to those shown here. Results are available upon request.

¹⁵ This specification, under homogenous preferences, disallows nonlinear effects of heat or cold, in which the WTP to avoid an additional day of extreme temperature increases with the number of days of extreme temperatures. It also disallows a preference for “seasonality,” which would be manifest as a lower WTP to avoid heat (cold) in a location with severe winters (summers). Nonlinear effects and preferences for seasonality are, however, captured by our heterogeneous preference specification, which allows the MWTP for any temperature bin to vary with the distribution of realized temperatures. We find that sorting and adaptation effects outweigh nonlinear damage effects, as discussed in section IV (especially for cold weather, see panels A and B of Figure 6). Overall, we do not find consistent evidence of a preference for seasons, as shown in Appendix Figure A1.1, which plots the MWTP for 40°F and 80°F as functions of CDD and HDD. These plots should slope upwards when households prefer seasons.

CDD, then $f(t)$ is a two-piece linear function with the kink point at 65°F. One of our objectives is to test whether $f(t)$ follows this functional form.

Given our data, the most flexible possible model for $f(t)$ would include a dummy variable for each 0.9°F temperature bin, in which each coefficient would signify a WTP relative to the bin containing 65°F. This model is impractical, however, as we have too little data to provide estimates that are precise enough to be meaningful.¹⁶ Instead, we model $f(t)$ using cubic splines per (5)

$$(5) \quad f(t) = \sum_{s=1}^S \beta_s S_s(t)$$

in which $S_1(t)$ through $S_S(t)$ are standard basis functions of a cubic B-spline of degree S . We space the knots of the basis functions evenly on the cdf of the population-weighted average temperature distribution.¹⁷ This spacing, rather than even spacing over the unweighted temperature support, clusters the nodes in the center of the distribution where the data are richest, improving flexibility in this region.

Defining N_{jt} as the average number of days per year at location j for which the average temperature is between t and $t + 0.9$, we flexibly estimate climate preferences by substituting into (4) the product of the spline function in (5) with N_{jt} , summed over all temperature bins:

$$(6) \quad \begin{aligned} \hat{Q}_j &= \sum_t N_{jt} f(t) + \beta_r \text{Rain}_j + \beta_h \text{Humid}_j + \beta_s \text{Sun}_j + \gamma \cdot \mathbf{D}_j + \xi_j \\ &= \sum_{s=1}^S \beta_s \left(\sum_t N_{jt} S_s(t) \right) + \beta_r \text{Rain}_j + \beta_h \text{Humid}_j + \beta_s \text{Sun}_j + \gamma \cdot \mathbf{D}_j + \xi_j \end{aligned}$$

¹⁶ We have experimented with wider bins and found that bins approximately 10°F wide are necessary for reasonable precision. In Table A1.1 in the Appendix, column R12 shows that such a binned specification yields similar qualitative results to those discussed below. We find the continuous splines more appealing because we do not believe that human comfort changes discontinuously every ten degrees.

¹⁷ The support of the contiguous U.S. temperature distribution ranges from -39.1°F to 111.2°F, covering 167 0.9°F bins. For a cubic spline of degree S , there are $S - 2$ nodes, including the nodes at the endpoints. In the 7th degree cubic spline that we ultimately focus on, the three interior nodes are located at 44.15, 58.55, and 71.15°F (these temperatures are the centers of the relevant bins).

As a reference case, we include in \mathbf{D}_j the “full set” of both geographic and demographic controls described above. In all of our regressions, we weight each PUMA (observation) by its population.¹⁸ For inference, we use standard errors that are clustered to allow for arbitrary spatial correlation of residuals across PUMAs within each state (Arellano 1987, Wooldridge 2003).¹⁹

B. Homogenous preference WTP estimates

Estimating the shape of $f(t)$

Figure 4 presents our estimates of $f(t)$ using 5th through 10th degree splines, given our reference case controls described in Section III.A. In each panel, the plotted estimate of the $f(t)$ curve is the MWTP for a day at temperature t rather than a day at 65°F, measured as a fraction of income. Each panel also depicts, for reference, the present and future (2070–2099, A2 scenario) U.S. population-weighted average temperature distributions. While the 5th degree spline appears too restrictive, and the 10th degree too noisy, several regularities emerge from the set of plots. First, the WTP curves consistently have an interior maximum near 65°F. We view this result—which is driven by the data and not “forced” by our QOL variable or functional form—as an important validation of our model and empirical strategy, since it accords with the intuition underlying HDD and CDD that WTP is maximized at 65°F. Second, there are too few days with average temperatures over 90°F to permit precise inference over this range, as evinced by the extremely wide standard error bands. Third, it appears that WTP departs non-linearly away from 65°F, undermining the restriction from the HDD/CDD model that $f(t)$ is linear. In particular, the slopes of the WTP curves tend to be steep on either side of 65°F but then level off at more extreme temperatures. Fourth, the decline in WTP away from 65°F is

¹⁸ As shown in Appendix Table A1.1, column R11, this weighting does not materially affect the estimates.

¹⁹ We have also experimented with clustering at the MSA level (in which PUMAs that are not part of an MSA are clustered within each state) and census division level. These approaches lead to estimated standard errors that are only slightly smaller and larger, respectively, than those presented here. When clustering on census divisions (of which there are nine) we use the cluster wild bootstrap to improve small sample performance (Cameron *et al.* 2008). This specification is displayed in column R14 of Appendix Table A1.1.

steeper on the right than on the left, meaning there is a greater WTP to avoid heat than to avoid cold.

The result that WTP declines less steeply over extreme temperatures accords with the intuition that households protect themselves from extremes by taking shelter in climate-controlled environments. Once the temperature is sufficiently uncomfortable that households spend little time outside, further increases in extreme temperature are less important. This result is consistent with research by Graff Zivin and Neidell (2012) that uses time-diary data to show that households spend less time outside in cold or hot weather.

The result that, on the margin, increases in heat are worse than increases in cold also follows intuition. Individuals can adapt to cold by wearing more clothing. However, options for thermoregulation are more limited in hot conditions.

The imprecision of our estimates at the extremes of the temperature distribution inhibit the ability of spline-estimated WTP functions to inform the welfare effects of climate change. This is true especially at the high end: days with average (not just high) temperatures exceeding 90°F are rare at present but common in climate change projections. To address this issue, we examine two functional forms that maintain flexibility in the interior of the temperature distribution while projecting WTP at the far extremes.²⁰ The first is a restricted 7th degree cubic spline, with WTP assumed to be constant over the extreme 1% of realized temperatures.²¹ The second is a four-piece linear spline, with kink points at 45°F, 65°F, and 80°F, allowing for a projection of decreasing WTP over extreme tem-

²⁰ We are not the first to confront the issue of conducting inference at temperatures near and beyond the limit of what is realized in present-day data. Prior work in the crop yield and health literatures has assumed that the damage function is constant beyond the point at which inference is no longer feasible (Deschênes and Greenstone 2007 and 2008, Schlenker and Roberts 2009). The first of our two restricted specifications—a 7th degree cubic spline with a restriction that MWTP is constant at the extremes—accords with this prior practice.

²¹ Specifically, we impose constant WTP over the bottom 0.5% and top 0.5% of the population-weighted temperature distribution. We apply this imposition after estimating the full 7th degree spline rather than beforehand, as reversing the order leads to unstable estimates at the cutoff points. We focus on a 7th degree spline because it visually appears to strike a balance between flexibility and precision. Using 6th or 8th degree splines instead does not substantially affect the results.

peratures.²² This second specification permits more straightforward hypothesis testing of how WTP changes over the temperature distribution.

Panels A and B of Figure 5, and columns I and II of Table 2, present results for both specifications, controlling for both geography and demographics. The point estimates in Table 2 can be interpreted as the WTP for each characteristic. In both specifications, WTP declines more steeply as temperatures increase away from 65°F than as they decrease from 65°F. With the linear spline, we reject a null hypothesis that the magnitudes of the slopes on either side of 65°F are equal with a p-value of 0.001. Both splines exhibit flatter slopes over the extremes than over the center of the distribution, particularly on the cold side of 65°F. For the linear spline, the change in slope at the 45°F kink point is statistically significant ($p = 0.005$), though the change in slope at 80°F is not ($p = 0.523$).

Thus, while we have some confidence in the result that extreme cold is not substantially more disamenable than moderate cold, the lack of statistical power over extremely hot days means that we are less confident in saying about extreme versus moderate heat.²³ The restricted cubic and linear spline specifications can therefore be viewed as two plausible alternatives. The former is conservative in assessing the WTP to avoid extreme heat because it assumes that WTP is constant over this range, whereas the linear spline model is more aggressive as it extrapolates the (negative) slope of the WTP function outside of the observed temperature range.

Non-temperature climate variables, controls, and robustness

The results in columns I and II of Table 2 indicate that households have a strong preference for sunshine, a mild preference for precipitation, and no statistically significant taste for humidity. Due to multi-collinearity, it is difficult to disentangle preferences for these three climate variables: locations that are very sunny also tend to be dry and non-humid. Amongst the geographic control variables, we estimate strong tastes for hilliness (average

²² The choices of 45°F and 80°F tend to yield the best fit to the data (per R^2) in most specifications. Alternative choices such as 50°F or 75°F do not substantially impact the preference or welfare estimates.

²³ The linear spline projection is sufficiently imprecise to the right of 80°F that we also cannot reject the hypothesis that WTP is constant over this range ($p = 0.117$).

land slope), a taste for proximity to the ocean, but no strong preference for proximity to a Great Lake. The estimated coefficients on the demographic control variables indicate that QOL increases with population density, educational attainment, average age, and percent Hispanic, while it decreases in percent black.

The demographic controls may themselves be endogenous, as demographic groups may have different tastes and sort accordingly.²⁴ These variables may also introduce an “overcontrolling” problem if they are, at least in part, determined by climate. That said, we believe that the benefits of controlling for demographics outweigh the costs. First, their inclusion substantially improves the precision of the estimated WTP for climate. Panels C and D in Figure 5, and columns III and IV of Table 2, present estimates of specifications that do not include the demographic variables. These specifications have notably wider WTP confidence intervals and less than one-half the R^2 of those that include demographics (using the cubic spline model, the standard errors on the estimated MWTPs for days at 40°F and 80°F are 26% and 55% larger, respectively, when the demographic controls are excluded). Second, including demographic variables helps guard against omitted variables bias and, in this case, provides evidence that such bias may not be a substantial concern. Although the demographic variables are powerful correlates of QOL, as evidenced by their large effect on R^2 , they have only a modest effect on the estimated WTPs for the climate variables. Nor, as we discuss in section V below, do they substantially impact our estimates of the welfare impact of climate change. We obtain these results despite evidence of some correlation of the demographic covariates—especially the percentage of high school graduates and the percentages of minorities—with warm

²⁴ The coefficient on population density is particularly difficult to interpret, since it captures unobserved amenities that come with population density (e.g., culture and restaurants) as well as idiosyncratic household-specific valuations of a location. The former drives the coefficient on population density upward, while the latter drives it downward (because the cost of living must fall (or wages must rise) to attract additional households). Empirically, the amenity aspects of population density appear to dominate, since we estimate a positive coefficient on this variable. Our estimated WTP for climate is robust to excluding population density from the specification. In section VI, we discuss the related issue of adjusting our QOL measure to account for changes in population between 1990 and 2000.

weather, as shown in Table A1.2.²⁵ Overall, following the logic of Altonji *et al.* (2005), our results suggest that if unobserved demographics substantially bias our climate preference and welfare estimates, they would have to be more strongly correlated with climate than the “headline” observable demographics included in our regressions.²⁶

Specifications that include state fixed effects (FE) along with demographics are examined in Figure 5 panels E and F, and columns V and VI of Table 2.²⁷ These specifications rely solely on within-state variation in climate for identification. Despite this substantial change in identifying variation, we still find in panel E that WTP is maximized near 65°F. The inclusion of state FE does result in a loss of precision and an increase in the estimated WTP to avoid cold. This increase may arise from the fact that the statistical power of the FE estimates comes from large, populous states with a large number of PUMAs and a broad range of climates. Thus, states like California and Texas are more influential than states in the Northeast and Midwest. If residents of these large, warm states have sorted because they are more averse to cold—a possibility consistent with our findings in section IV below—this may explain the difference between the FE and non-FE estimators. Nonetheless, the estimated WTP to avoid cold remains smaller than the estimated WTP to avoid heat, though this difference is no longer statistically significant. The estimates are

²⁵ As an additional robustness check, we have found that adding squared terms of the demographic variables does not substantially affect the results. See Table A1.1, column R9. It may be that the semiparametric temperature functions we use are less susceptible to the raw correlations shown in Table A1.2, or that the partial correlations after controlling for the other weather variables and geographic controls are relatively small. We have also controlled for public spending on parks (though data are available for only 1,970 PUMAs) and found that these controls do not qualitatively affect our estimates.

²⁶ In other settings for hedonic studies—for example, measurements of the WTP for clean air—it has been shown that estimates are sensitive to the inclusion or exclusion of demographic controls (see, for example, Chay and Greenstone 2005). In such cases, the source of bias is often intuitive. For example, emitters of pollutants are often concentrated in industrial areas that tend to be populated by lower-income households. With regard to climate, however, a source of omitted variable bias from demographics is less obvious *ex-ante* because a wide range of demographic groups and urban vs. rural locations can be found in essentially every climate zone. Our result that adding demographic variables to the specification does not substantially affect the estimated WTP for climate accords with this intuition.

²⁷ As shown in Appendix Table A1.1, specification R10, using census division fixed effects rather than state fixed effects does not result in a substantial difference relative to the non-FE estimate.

sufficiently imprecise that the state FE estimates are not statistically significantly different than those without FE.²⁸

Estimated preferences derived from additional alternative specifications are presented in Appendix Table A1.1. These regressions explore alternative control variables, specifications for $f(t)$, subsamples of the data (e.g., omitting superstar cities), and weighting schemes. This table also provides estimated impacts of climate change on amenity values and estimates that allow for preference heterogeneity, as discussed below. Overall, our conclusions do not qualitatively change across these specifications.

IV. Estimation of WTP for Climate with Heterogeneous Preferences

A. Empirical strategy

Section III above and prior work on climate amenity valuation has assumed that households share homogenous preferences for climate. Here, we relax this assumption and allow households to sort into locations that suit their preferences (noting that heterogeneity may also reflect households' adaptation to their local climate). Estimation under these more relaxed conditions is based on the framework developed by Bajari and Benkard (2005) and applied by Bajari and Kahn (2005). The intuition of this approach lies in the first order condition of equation (3) and is illustrated in Figure 1, which depicts an equilibrium in which the only characteristic is average summer temperature, T_s . Given a nonlinear hedonic price function $\hat{Q}(\mathbf{Z}_j, \xi_j)$, the MWTP of households located at j for a characteristic k (e.g., a day in a given temperature bin) is simply given by $\partial \hat{Q}(\mathbf{Z}_j, \xi_j) / \partial Z_k$. Thus, flexible estimation of $\hat{Q}(\mathbf{Z}_j, \xi_j)$ allows us to recover the distri-

²⁸ We test the null hypothesis that the state FE are uncorrelated with the climate variables via block wild bootstrapping, clustering on state (200 repetitions). For each draw, we estimate preferences both with and without state FE. We then use all draws to obtain a standard deviation of the difference between the estimators. This procedure, unlike a traditional Hausman test, does not require that the non-FE estimator be efficient (see Cameron *et al.* 2005, p. 718). We find that we cannot reject equality of the WTP for 40°F across the estimators (the p-values are 0.858 for the cubic spline and 0.603 for the linear spline) nor equality of the WTP for 80°F (the p-values are 0.445 for the cubic spline and 0.560 for the linear spline). We do, however, marginally reject equality of the welfare impact of climate change across the two estimators (the p-values are 0.086 for the cubic spline and 0.076 for the linear spline).

bution of MWTP for each characteristic k across the population of households. In Figure 1, for example, the depicted equilibrium is consistent with positive sorting in which households with a high MWTP to avoid heat settle in areas with low summer temperatures. In contrast, a concave equilibrium price function would reflect non-linear damages from heat, so that the MWTP to avoid heat would be greater for hot locations than for mild locations.

To flexibly estimate $\hat{Q}(\mathbf{Z}_j, \xi_j)$, we use local linear regression per Fan and Gijbels (1996). We suppose that, local to location j^* , $\hat{Q}(\mathbf{Z}_j, \xi_j)$ satisfies (7) below:

$$(7) \quad \hat{Q}_j = \sum_k \beta_k^{j^*} Z_{jk} + \xi_j$$

In (7), the implicit prices β are superscripted by j^* , as we estimate a distinct set of prices at each location. We obtain the β^{j^*} at each location via weighted least squares per (8) and (9):

$$(8) \quad \beta^{j^*} = \arg \min_{\beta} (\mathbf{Q} - \mathbf{Z}\beta)' \mathbf{W} (\mathbf{Q} - \mathbf{Z}\beta)$$

$$(9) \quad \mathbf{Q} = [\mathcal{Q}_j]; \quad \mathbf{Z} = [\mathbf{Z}_j]; \quad \mathbf{W} = \text{diag} [K_h(\mathbf{Z}_j - \mathbf{Z}_{j^*})]$$

where \mathbf{W} is a matrix of kernel weights defined so that locations that are similar to j^* in characteristics receive the most weight in the regression.²⁹ This approach allows households in relatively hot or cold locations to have different MWTPs to avoid departures from mild climates. To calculate \mathbf{W} , we use a normal kernel function with a bandwidth h

²⁹ In principle, \mathbf{W} can include every characteristic in the model, so that preferences can vary in the cross-section not only with the climate variables but also with the controls. In practice, however, we include only the temperature spline basis functions in \mathbf{W} , since including the controls requires a very large bandwidth in order to avoid collinearity in the weighted covariate matrix. In addition, when implementing our procedure with the restricted 7th degree spline specifications, we only include a 5th degree cubic spline in \mathbf{W} . Including the full 7th degree cubic spline in \mathbf{W} results in very noisy estimates of temperature preferences at extreme locations (such as northwest Minnesota) that have few neighbors in climate space.

of 2, per (10) and (11) below.³⁰ $\hat{\sigma}_k$ denotes the standard deviation of characteristic k across the sample:

$$(10) \quad K(\mathbf{Z}) = \prod_k N(Z_k / \hat{\sigma}_k)$$

$$(11) \quad K_h(\mathbf{Z}) = K(\mathbf{Z} / h) / h$$

Ideally, our estimator would allow for heterogeneity in households' MWTP for all local characteristics, not just the temperature profiles. Were the choice set continuous in characteristics, point identification of preferences for all characteristics at each location would indeed be possible. In our setting, however, there is a discrete number of 2,057 PUMAs, which implies that households' preferences are only set identified. That is, there is a range of MWTPs such that each household would prefer to stay in its PUMA rather than move. We study these identified sets using a Gibbs sampling method per Bajari and Benkard (2005) and find that they are often large enough to be uninformative when households are permitted to have heterogeneous preferences over all characteristics (including the control variables).³¹ Our main specifications therefore maintain the restriction that households share a homogenous MWTP for non-temperature characteristics.³² With this restriction, we find that the identified sets of MWTPs for temperature are sufficiently small to closely approximate point identification. The results presented below therefore follow Bajari and Kahn (2005) in treating MWTP as point identified at each location.

³⁰ We follow Bajari and Benkard (2005) and Bajari and Kahn (2005) in choosing a normal kernel and in choosing a bandwidth that yields visually appealing preference estimates, in the sense that estimated MWTP varies smoothly over characteristics space (as in Figure 6). In contrast, using leave-one-out cross-validation we find an "optimal" bandwidth of 0.22 in the cubic spline preference specification. This small bandwidth leads to very noisy and imprecise preference estimates as well as an extremely large estimated amenity loss from climate change (18.5% of income under the A2 scenario). In general, the estimated amenity loss increases as the bandwidth becomes smaller, so our bandwidth choice of 2 is conservative.

³¹ Bajari and Benkard (2005) find similar results when the number of parameters over which preferences may be heterogeneous is permitted to increase, relative to a fixed choice set.

³² Prior to running the local linear regression, we enforce preference homogeneity over the non-temperature variables by stripping their effects from \hat{Q} using the OLS estimates from section III (this approach follows an example given in Bajari and Benkard 2005). Thus, the only characteristics involved in the local linear regression are the temperature spline basis functions.

Appendix 4 presents set identified estimates and discusses our set identification procedure in detail.

Finally, we emphasize that we only identify the MWTP of each household for each temperature at its chosen climate and cannot identify its WTP for a large change in climate, over which MWTP may vary. In Figure 1, for instance, the MWTPs at locations A and B are identified, but the shapes of the indifference curves are not. Estimating the welfare impacts from non-marginal changes in climate therefore requires assuming a functional form for households' utility. We focus on welfare estimates that assume the indifference curves are linear, with a slope equal to the estimated MWTP. This approach is transparent, permits an apples-to-apples comparison to the homogenous preference estimates, and is conservative: allowing for concavity would unambiguously increase the estimated amenity losses, as the cost of each additional hot day would rise (and the benefit from each reduction in cold days would fall).

The assumption that households' utility functions are linear in characteristics creates the problem, however, that the second order condition for utility maximization may not hold if the hedonic price function is not globally convex. Because we ultimately find that the $\hat{Q}(Z_j, \xi_j)$ function is slightly concave in some dimensions, we have studied a specification in which households' utility functions are sufficiently concave that their utility is maximized at their present location. Our procedure for estimating this specification is discussed in appendix 3. We find that allowing for concavity in the utility function increases the estimated losses from climate change by roughly 0.4 percentage points relative to the results given in section V below (for the cubic spline model in column I of tables 3 and 4).

B. Heterogeneous preference MWTP estimates

Estimated MWTPs for temperature by PUMA are graphed in Figure 6: panels A and C from the restricted cubic spline, panels B and D from the linear spline, all controlling for geography and demography, but not state FE. Panels A and B present the distribution of MWTP, scaled up by 365, for an additional day at 40°F relative to a day at 65°F, as a

function of HDD. Each of the plotted points represents the preferences of a specific PU-MA in the data. The strong upward slope in both panels indicates that households with the most negative MWTP for cold weather tend to be located in areas with the fewest HDDs (the reduction in the MWTPs' magnitudes near zero HDD or CDD may simply reflect large confidence intervals; see standard deviations for the MWTP estimates in appendix Figure A1.6). This result is consistent both with households sorting based on their aversion to cold and with adaptation to cold climates. Panels C and D, in contrast, examine the MWTP for an additional day at 80°F as a function of CDD. Panel C, which uses the restricted cubic spline model, exhibits a weak downward slope, consistent with residents of hot areas being more heat-averse on the margin and indicating mild concavity in households' utility functions rather than sorting. This result also suggests that households have only a limited ability to adapt to heat, relative to their ability to adapt to cold. The linear spline model in Panel D, however, indicates a positive slope, suggestive of some sorting on willingness to avoid heat, though the overall heterogeneity in MWTP is still smaller than that for the willingness to avoid cold.

An alternative depiction of heterogeneous climate preferences is given in Appendix Figure A1.2. This figure plots MWTP curves under the cubic spline model at a few select cities. Beyond showing that the magnitude of households' MWTP to avoid cold is negatively correlated with the amount of cold they face, this figure also shows that the estimated maximum of the MWTP function occurs at slightly higher temperatures in warmer locations, again consistent with sorting based on aversion to cold or with adaptation.

The average MWTP to avoid 40°F and 80°F for the cubic and linear spline models plotted in Figure 6 are given in columns I and II of Table 3. The average MWTP to avoid heat exceeds that to avoid cold, consistent with the homogenous preference model (columns I and II of Table 2). Columns III through VI of Table 3 proceed through the same control variable specifications that were used in Table 2, with similar results. Appendix Table A1.1, panel (b) provides estimates from numerous alternative specifications.

V. Estimated Impacts of Climate Change on Climate Amenities

In this section, we use climate preference estimates to calculate how predicted climate change may impact local amenity values. These results maintain technology and preferences (specifically, WTP as a share of income) at their present levels and are best interpreted as welfare impacts that would occur if changes predicted for the end of the century were to occur now, allowing for spatial adaptation (e.g., households in the North adopting southern behaviors and levels of air conditioning penetration) to occur instantly.

We focus first on the reference case model I from Table 2, using the restricted cubic spline under homogenous preferences, with projections from the A2 scenario for 2070–2099. Table 4 displays estimated changes in welfare for this specification (column I) and several variants. We project a net welfare loss of 2.28% of personal income, or \$195 billion per year (in year 2000 dollars),³³ from decreased amenity value of climate. This loss is driven by the result that the WTP to avoid heat exceeds the WTP to avoid cold, so that the welfare loss from hotter summers of -4.85% exceeds the gain from warmer winters of 2.12%. These temperature-driven welfare losses are slightly offset by increases in amenities from the small predicted changes to precipitation, humidity, and sunshine. Results from the linear spline in Column II predict slightly higher welfare costs, with higher benefits from warmer winters and greater costs from hotter summers. Finally, the bottom row of Table 4 presents welfare impacts under the alternative A1FI scenario, which predicts more extreme changes in temperature. This scenario results in larger welfare effects from milder winters and hotter summers, so that net welfare losses are roughly 3% of total income.

The geographic distributions of welfare changes for the restricted cubic and linear spline specifications are shown in panels A and B of Figure 7. Both maps show a few areas in the Pacific Northwest, northern Michigan, and the upper Northeast—accounting for about 6% of the population—for which we forecast a minor improvement. The rest of the U.S. population will experience losses. These are especially severe in the Southeast and

³³ The dollar figure is based on year 2000 total personal income of \$8.576 trillion in the 48 contiguous states.

Southwest: some areas in Florida, Louisiana, Texas and California have expected losses greater than 4 percent of income. Much of the country, as far north as the Dakotas, is expected to experience substantial losses as they lose many days with moderate temperatures, which we estimate are highly valued.

The predictions of the two estimated temperature functions agree less in the center of the country, where the restricted cubic spline predicts relatively small losses. This result occurs because summer temperatures in this area are already sufficiently hot that residents are on the flat portion of the WTP curve. In contrast, the linear spline predicts relatively large welfare losses, as already hot summer days become hotter. These predictions are imprecise as the historical climate record offers too few extremely hot days to allow a precise claim about WTP over such extreme heat.

Appendix Table A1.1 presents estimates of amenity losses from several other alternative specifications with varying control structures, generally finding similar results, with welfare losses of 2 percent or more of income. Our results are largely unchanged if we remove the demographic controls, although they do become less precise. The greatest changes occur if we use estimates that include state FE, as seen in column III of Table 4. These estimates imply greater welfare gains from warmer winters and smaller losses from hotter summers, resulting in a smaller net welfare loss closer to one percent of income.

Accounting for heterogeneous preferences yields a somewhat larger net amenity loss of 2.79% of income, as shown in column IV of Table 4. This difference is primarily driven by the finding that households in the North are the least averse to cold. As a result, the predicted reduction in cold days is negatively correlated spatially with the MWTP to avoid those cold days, so that those who see the greatest reduction of winter are those who would benefit from it the least. The spatial distributions of estimated welfare losses for specifications IV and V in Table 4 are shown in appendix Figure A1.7

For small changes in climate, the estimated welfare effects shown in Table 4 provide accurate first-order approximations, robust to migration via the envelope theorem. However, if climate changes are sufficiently large, the second-order effect of migration may be material. We assess the importance of migration with a parsimonious model under

homogenous preferences. Intuitively, households will leave areas with worsening climates and migrate towards areas with improving climates. We model the case in which mobility responses are proportional to changes in QOL, $\Delta\hat{Q}_j$, and households remain within the United States. Thus, changes in the logarithm of population density of location j , $\Delta\hat{N}_j$, depend on changes in QOL at j relative to the average QOL change in the United States, $\Delta\hat{Q}_{avg}$, per the formula

$$(12) \quad \Delta\hat{N}_j = \varepsilon_{N,Q}(\Delta\hat{Q}_j - \Delta\hat{Q}_{avg})$$

where $\varepsilon_{N,Q}$ is the long-run elasticity of population with respect to changes in quality of life. Although direct empirical evidence on this elasticity is scant, work by Albouy and Stuart (2014) suggests a value of 8.³⁴ In other words, a one percent increase in QOL will lead to an eight percent increase in population density, taking into account long-run responses in housing and migration.

Allowing for adaptive migration per equation (12) we estimate that for our baseline specification I, the average absolute value of the population change across all PUMAs is 10.3%. This migration only modestly changes our estimated aggregate welfare loss from 2.28% to 2.01%. Adaptive migration only slightly mitigates welfare losses because only a few areas actually improve, and migration to these areas is constrained by crowding.

VI. Alternative measures of QOL

In section I above, we emphasized the importance of correctly weighting wage and housing-cost and wage differentials when calculating our QOL measure, while also making a correction for commuting costs. In Table 5, we examine climate-preference and welfare estimates based on alternative measures. First, in column II, we use wage differ-

³⁴ This elasticity incorporates five different responses to a change in QOL: i) housing producer substitution away from land; i.e. housing supply elasticity ii) traded-good producers' substitution away from land, creating diminishing returns from fixed factors, iii) consumers' substitution away from housing; i.e. crowding into the existing housing stock, iv) willingness of households to consume less overall (net of housing) in high QOL areas; and v) endogenous agglomeration economies and congestion effects from changes in population density.

entials by place of residence rather than place of work, without correcting for commuting. This change slightly attenuates the WTP estimates, leading to smaller inferred damages. Intuitively, the wage and commuting correction to QOL does not substantially impact the WTP estimates because these factors operate on a small geographic scale that is likely orthogonal to climate variation. We next consider the separate roles that wages and housing costs play in determining QOL. In column III, we measure QOL using only the additive inverse of wages (so that low wages denote higher QOL). Without state FE, this change causes extreme temperatures to be valued positively; with fixed effects, the specification finds a positive (though statistically insignificant) WTP for heat, recalling the results of Hoch and Drake (1974) and Moore (1998).³⁵ Conversely, when only housing costs are used to measure QOL in column IV, the estimates exaggerate households' aversion to extreme temperatures and damages to climate change. Column V uses a QOL measure similar to Roback (1982), Nordhaus (1996), and Sinha and Cropper (2013) that, relative to our measure, puts 2.7 times more weight on wages relative to housing costs (by not accounting for taxation and non-housing local costs). Not surprisingly, these results resemble those based only on wages, albeit less extremely. The dramatic difference in results between the specifications with and without fixed effects also evokes the lack of robustness that Nordhaus (1996) finds using his QOL measure.

Columns VI and VII of Table 5 return to our baseline QOL measures but alternatively use wage levels (and income and expenditure shares s_w and s_y) for workers without a college degree or with one, to test whether these two groups value amenities differently. Our findings suggest that more educated workers may be relatively heat-averse.

As discussed in section I, a potential concern regarding our baseline QOL measures is that our estimates of the PUMA-level wage differentials \hat{w}_j are not robust to sorting based on unobserved skills or a spatial match component of wages. We address this concern in several ways. First, column VIII of Table 5 presents results from a specification in

³⁵ In the column III state FE specification, the estimated welfare impact is similar to that obtained with the reference case QOL measure despite the very different MWTPs at 40°F and 80°F. This result occurs only because the estimated MWTP is very large and positive over extreme cold <30°F. The estimates in column V with state FE behave similarly.

which we adopt the procedure used in Dahl (2002) to adjust the \hat{w}_j for selective migration. In our wage equation, we include a flexible control function for each individual's probability of moving from his or her state of birth to the current state, as well as the probability of staying in the state of birth, based on the individual's observed demographic characteristics.³⁶ This adjustment will decrease the estimated \hat{w}_j (and therefore increase the estimated \hat{Q}_j) for PUMAs with a relatively high rate of in-migration to the extent that in-migrants earn relatively high wages. As shown in column VIII, this adjustment to our QOL estimates does not substantially affect our preference and welfare estimates.

As an alternative and closely related approach for addressing selective migration, we have also adjusted our baseline QOL estimates for each PUMA by adding a correction equal to the PUMA's rate of net in-migration between 1990 and 2000 (in percent) multiplied by a mobility cost estimate from Notowidigdo (2012).³⁷ This adjustment explicitly increases the estimated QOL for locations that are growing relatively quickly. We find that this adjustment modestly increases the estimated aversion to cold and slightly attenuates the estimated welfare loss under the A2 model (by 0.3 percent of income). Finally, we have estimated a specification that insulates our estimates from selective migration to "superstar cities" by dropping superstar metropolitan areas (as defined in Gyourko, Mayer, and Sinai (2013)) from the sample (we drop 442 of the 2,057 PUMAs in the sample).

³⁶ Per Dahl (2002), we estimate these probabilities by category according to various demographic characteristics. We use the exact same categories as Dahl (2002) for movers (20 for each state) and stayers (70). We also add a separate mover category / "birth-state" for those born outside the U.S. As Dahl only used male, white, and ages 25-34, we create twelve times the number of original categories to account for female, non-white, and age categories 35-44 and 45-55. To identify the wage intercepts across states, we constrain the coefficients in the control function to be the same across states, unlike in Dahl (2002).

³⁷ Notowidigdo (2012) allows for heterogeneity in mobility costs and estimates a mobility cost elasticity, which dictates how the mobility cost of the marginal migrant increases with the number of migrants. This elasticity is appropriate for adjusting our QOL estimates because it is the marginal migrant that determines market prices and wages. Our calculations use a weighted average of Notowidigdo's (2012) estimated elasticities for high-skill and low-skill workers (0.066 and 0.065 respectively, see Table 7 in his paper). These estimates are identified from plausibly exogenous Bartik (1991) labor demand shocks and correspond to migration over a ten-year period, thereby aligning with our 1990 to 2000 migration rates. We also use Notowidigdo's (2012) curvature parameters (betas) in specifying the mobility cost function.

We find that doing so also modestly decreases the estimated MWTPs and welfare impact, without qualitatively affecting our conclusions (see Appendix Table A1.1, column R15).

VII. Conclusions

This paper illustrates the ability of the Rosen-Roback hedonic framework, when properly calibrated, to recover willingness to pay for non-market goods such as climate. By using a quality of life measure that is carefully constructed from local wage and housing price differentials in the contiguous U.S. and by adopting a flexible model of climate based on local temperature distributions, we obtain estimates of households' valuation of climate amenities that both accord with intuition and are generally robust to a wide range of empirical specifications. We find that households: (1) place the most value on temperatures near 65°F, (2) tend to dislike marginal increases in heat more than marginal increases in cold, and (3) suffer less from marginal increases in cold or heat once the temperature is sufficiently cold or hot (though inference is weak for hot weather). Finally, we find evidence of heterogeneity in these preferences, with households that are most averse to cold living in the South, consistent with models of both sorting and adaptation.

When we apply our estimates to value the amenity impacts of climate change, we estimate an average U.S. welfare loss between one and four percent of income per year by 2070–2099 under “business as usual.” Most areas in the United States are predicted to experience losses, as many pleasantly warm days will be lost while uncomfortably cold days will only be moderately abated. These estimates are similar in magnitude to the percentage losses in GDP (market goods) predicted by Nordhaus (2007) and the Stern report (2007) for 2100 under business as usual, suggesting that nonmarket amenity benefits alone may justify the mitigation costs discussed in their studies.³⁸

The greatest limitation to our empirical method is that we have little statistical power to make inferences regarding households' willingness to pay (WTP) to avoid temperatures beyond the limits of the current temperature distribution. The data provide only limited

³⁸ The Nordhaus and Stern market damage estimates of 3% and 2.9% apply to worldwide GDP, not U.S. GDP.

guidance as to how WTP should be extrapolated over extremely hot temperatures that have yet to be realized. Should the WTP to avoid heat actually increase non-linearly over the extremes, our damage estimates will be too conservative, particularly in the southern U.S.

References

- Albouy, David Y. (2011) "Are Big Cities Bad Places to Live? Estimating Quality of Life across Metropolitan Areas." NBER Working Paper #14472 (original 2008, revised December 5, 2011).
- Albouy, David Y. and Bert Lue (forthcoming) "Driving to Opportunity: Local Wages, Commuting, and Sub-Metropolitan Quality of Life," *Journal of Urban Economics*.
- Albouy, David Y. and Bryan Stuart (2014) "Urban Population and Amenities," NBER Working Paper 19919.
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber (2005) "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of Political Economy* 113, pp. 151-184.
- Arellano, Manuel, "Computing Robust Standard Errors for Within-Groups Estimators," *Oxford Bulletin of Economics and Statistics*, 49 (1987), pp. 431-434.
- Bajari, Patrick and C. Lanier Benkard (2005) "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach," *Journal of Political Economy* 113, pp. 1239-1276.
- Bajari, Patrick and Matthew E. Kahn (2005) "Estimating Housing Demand with an Application to Explaining Racial Segregation in Cities," *Journal of Business and Economic Statistics* 23, pp. 20-33.
- Barreca, Alan (2012) "Climate Change, Humidity, and Mortality in the United States," *Journal of Environmental Economics and Management* 63, pp. 19-34.
- Barreca, Alan, Karen Clay, Olivier Deschênes, Michael Greenstone, and Joseph S. Shapiro (2013) "Adapting to Climate Change: The Remarkable Decline in the U.S. Temperature-Mortality Relationship over the 20th Century," NBER working paper

#18700.

- Bartik, Timothy J. (1991) "Who Benefits from State and Local Economic Development Policies," W.E. Upjohn Institute for Employment Research: Kalamazoo, MI.
- Baum-Snow, Nathaniel and Ronni Pavan (2012) "Understanding the City Size Wage Gap," *Review of Economic Studies* 79, pp. 88-127.
- Blomquist, Glenn C, Mark C. Berger and John P. Hoehn (1988) "New Estimates of Quality of Life in Urban Areas." *American Economic Review* 78, pp. 89-107.
- Cameron, A. Colin and Pravin K. Trivedi (2005) *Microeconometrics: Methods and Applications*, New York: Cambridge University Press.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller (2008) "Bootstrap-Based Improvements for Inferences with Clustered Errors," *Review of Economics and Statistics* 90, pp. 414-427.
- Chay, Kenneth Y. and Michael Greenstone (2005), "Does Air Quality Matter? Evidence from the Housing Market," *Journal of Political Economy* 113, pp. 376-424.
- Chen, Yong and Stuart S. Rosenthal (2008) "Local Amenities and Life-Cycle Migration: Do People Move for Jobs or Fun?" *Journal of Urban Economics* 64, pp. 519-537.
- Cragg, Michael I. and Matthew E. Kahn (1997) "New estimates of climate demand: evidence from location choice." *Journal of Urban Economics* 42, pp. 261-284.
- Cragg, Michael I. and Matthew E. Kahn (1999) "Climate consumption and climate pricing from 1940 to 1990," *Regional Science and Urban Economics* 29, pp. 519- 539.
- Dahl, Gordon B. (2002) "Mobility and the Return to Education: Testing a Roy Model with Multiple Markets," *Econometrica* 70, pp. 2367-2420.
- Deschênes, Olivier and Michael Greenstone (2007) "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather," *American Economic Review* 97, pp. 354-385.
- Deschênes, Olivier and Michael Greenstone (2011): "Climate Change, Mortality and Adaptation: Evidence from Annual Fluctuations in Weather in the U.S.," *American Economic Journal: Applied Economics* 3, pp. 152-185.
- Deschênes, Olivier and Michael Greenstone (2012) "The Economic Impacts of Climate

- Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Reply,” *American Economic Review*, 102, pp. 3761-3773.
- Fan, Jianqing and Irène Gijbels (1996), *Local Polynomial Modelling and Its Applications*. London: Chapman and Hall.
- Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker (2012) “The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment,” *American Economic Review* 102, pp. 3749-3760.
- Frijters, P. and Van Praag, B.M.S. (1998) “The effects of climate on welfare and wellbeing in Russia.” *Climatic Change* 39, pp. 61–81.
- Glaeser, Edward L. and David C. Maré (2001) “Cities and Skills,” *Journal of Labor Economics* 19, pp. 316-342.
- Graff Zivin, Joshua and Matthew Neidell (2012) “Temperature and the Allocation of Time: Implications for Climate Change,” *Journal of Labor Economics*, forthcoming.
- Graves, Philip E. (1979) “A Life-Cycle Empirical Analysis of Migration and Climate, by Race.” *Journal of Urban Economics* 6, pp. 135-47.
- Gyourko, Joseph, Christopher Mayer, and Todd Sinai (2013) “Superstar Cities,” *American Economic Journal: Economic Policy* 5, pp. 167-199.
- Hoch, Irving and Judith Drake (1974) “Wages, Climate, and the Quality of Life.” *Journal of Environmental Economics and Management* 1: pp. 268-296.
- International Panel on Climate Change Working Group I (2007) *Climate Change 2007: The Physical Science Basis*. Published for the International Panel on Climate Change.
- Kahn, Matthew E. (2009) “Urban Growth and Climate Change.” *Annual Review of Resource Economics* 1, pp. 333-350.
- Kennan, John and James R. Walker (2011), “The Effect of Expected Income on Individual Migration Decisions,” *Econometrica* 79, pp. 211-251.
- Kuminoff, Nicolai V. and Jaren C. Pope (2012) “Do ‘Capitalization Effects’ for Public Goods Reveal the Public’s Willingness to Pay?” working paper.
- Maddison, David (2003) “The Amenity Value of the Climate: the Household Production

- Function Approach,” *Resource and Energy Economics* 25, pp. 155-175.
- Maddison, David and Bigano, A. (2003) “The Amenity Value of the Italian Climate.” *Journal of Environmental Economics and Management* 45, pp. 319-332.
- Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw (1994) “The Impact of Global Warming on Agriculture: A Ricardian Analysis,” *American Economic Review* 84, pp. 753-771.
- Moore, Thomas Gale (1998) “Health and Amenity Effects of Global Warming.” *Economic Inquiry* 36, pp. 471–488.
- Nordhaus, William D. (1996) “Climate amenities and global warming”, in *Climate Change: Integrating Science, Economics, and Policy*. N. Nakicenovic, W.D. Nordhaus, R. Richels, and F. Toth, eds. IIASA, pp. 3–45.
- Nordhaus, William D. (2007) *The Challenge of Global Warming: Economic Models and Environmental Policy*. Vol 4. Yale University.
- Notowidigdo, Matthew J. (2012), “The Incidence of Local Labor Demand Shocks,” mimeo.
- Peiser, Richard B. and Lawrence B. Smith (1985) “Homeownership Returns, Tenure Choice and Inflation,” *American Real Estate and Urban Economics Journal* 13, pp. 343-360.
- Rehdanz, Katrin (2006) “Hedonic Pricing of Climate Change Impacts to Households in Great Britain.” *Climatic Change*, 74, pp. 413-434.
- Rehdanz, Katrin, and David J. Maddison (2005) “Climate and Happiness.” *Ecological Economics* 52, pp. 111–25.
- Rehdanz, Katrin and David J. Maddison (2009) “The Amenity Value of Climate to German Households.” *Oxford Economic Papers* 61, pp. 150-167.
- Roback, Jennifer (1982) "Wages, Rents, and the Quality of Life." *Journal of Political Economy* 90, pp. 1257-1278.
- Rosen, Sherwin (1974) “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition.” *Journal of Political Economy* 82, pp. 34-55.
- Rosen, Sherwin (1979) "Wages-based Indexes of Urban Quality of Life," in *Current Is-*

- sues in Urban Economics*. P. Mieszkowski and M. Straszheim, eds., Baltimore: John Hopkins Univ. Press.
- Savaeau, David (1999) *Places Rated Almanac*. Foster City, CA: IDG Books Worldwide.
- Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher (2006) “The Impact of Global Warming on U.S. Agriculture: An Econometric Analysis of Optimal Growing Conditions.” *Review of Economics and Statistics* 88, pp. 113-125.
- Schlenker, Wolfram, and Michael J. Roberts (2009) “Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields Under Climate Change.” *Proceedings of the National Academy of Sciences*, 106, pp. 15594-15598.
- Shapiro, P. and Smith, T., (1981) “Preferences on Non Market Goods Revealed Through Market Demands.” *Advances in Applied Microeconomics: A Research Annual* 1, pp. 105-122.
- Sinha, Paramita and Maureen L. Cropper (2013) “The Value of Climate Amenities: Evidence from U.S. Migration Decisions.” NBER Working Paper 18756.
- Stern, Nicholas (2007) *The Economics of Climate Change: the Stern Review*. Cambridge University Press.
- Timmins, Christopher (2007) “If You Can’t Take the Heat, Get Out of the Cerrado...Recovering the Equilibrium Amenity Cost of Nonmarginal Climate Change in Brazil.” *Journal of Regional Science* 47, pp. 1-25.
- Tol, Richard S. J. (2002) “Estimates of the Damage Costs of Climate Change, Part II. Dynamic Estimates.” *Environmental and Resource Economics* 21, pp. 135–160.
- Tol, Richard S. J. (2009) “The Economic Effects of Climate Change.” *Journal of Economic Perspectives* 23, pp. 29–51.
- Wooldridge, Jeffrey M., “Cluster-Sample Methods in Applied Econometrics,” *American Economic Review Papers and Proceedings*, 93 (2003), 133-138.

Figure 1: Illustrative hedonic price function with demand-side equilibrium FOC's satisfied

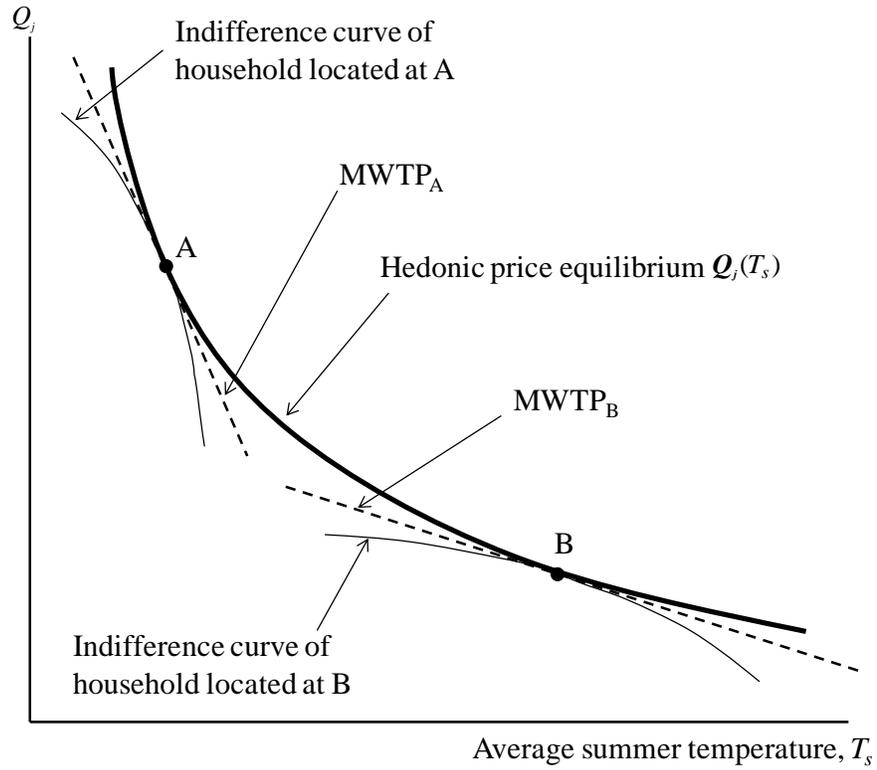


Figure 2: Quality of life (QOL) differentials in 2000

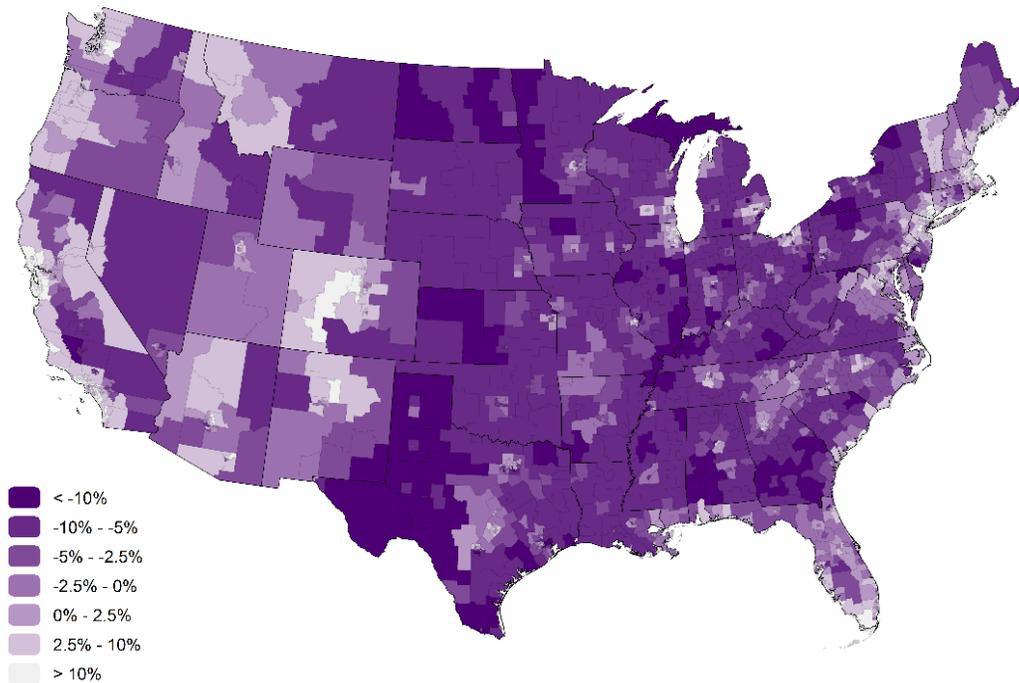
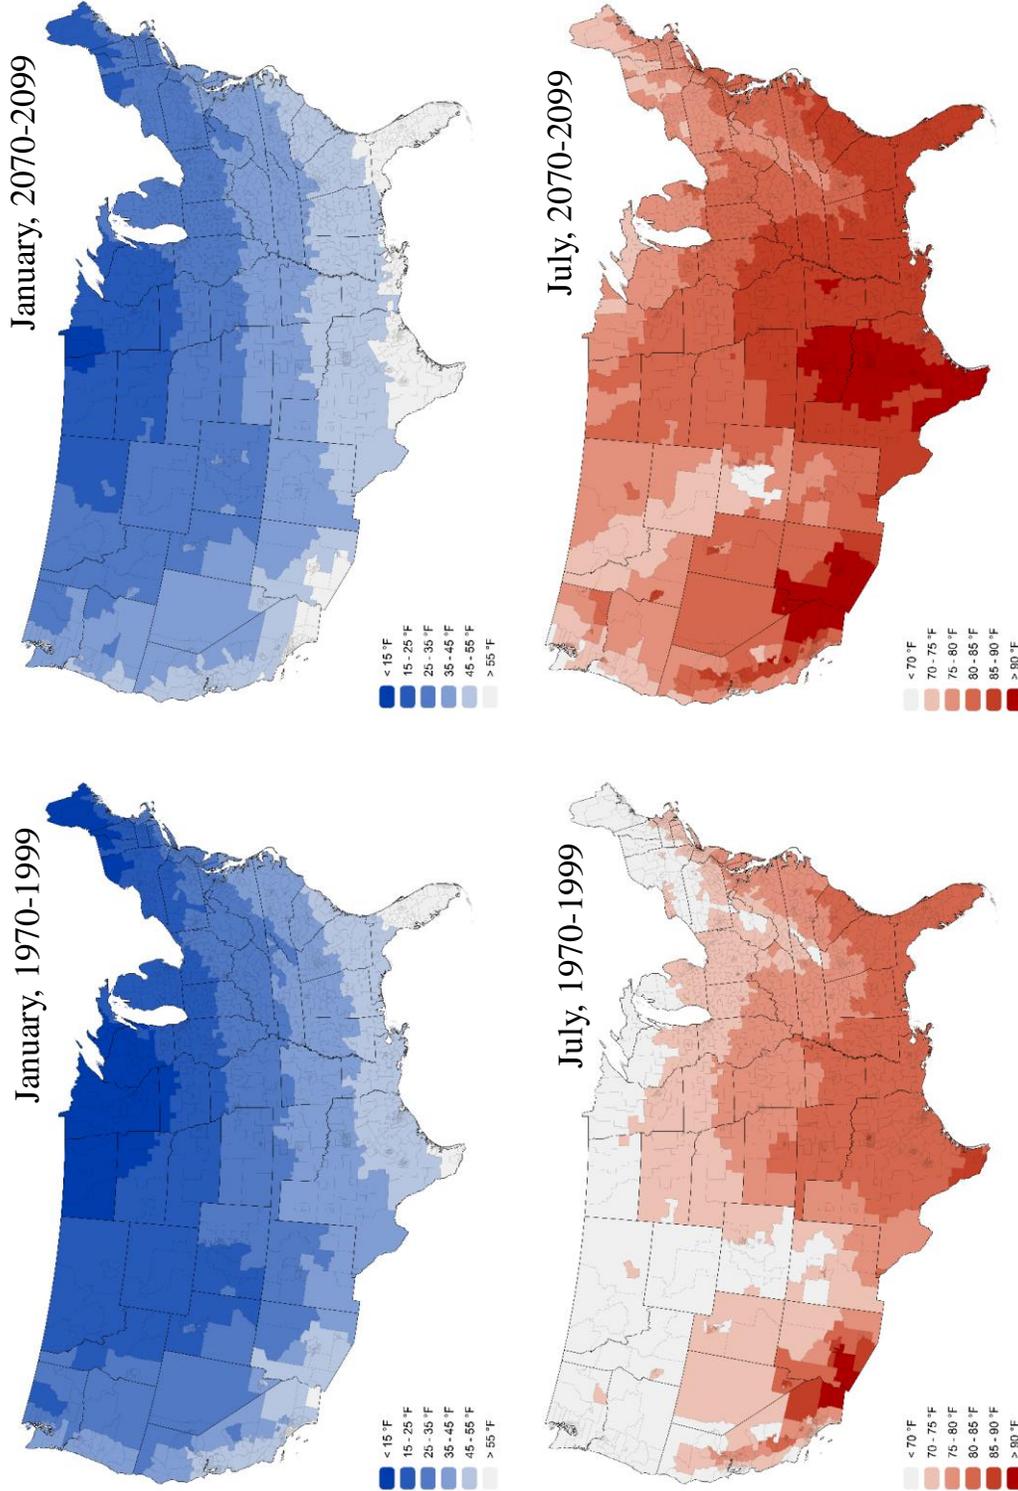
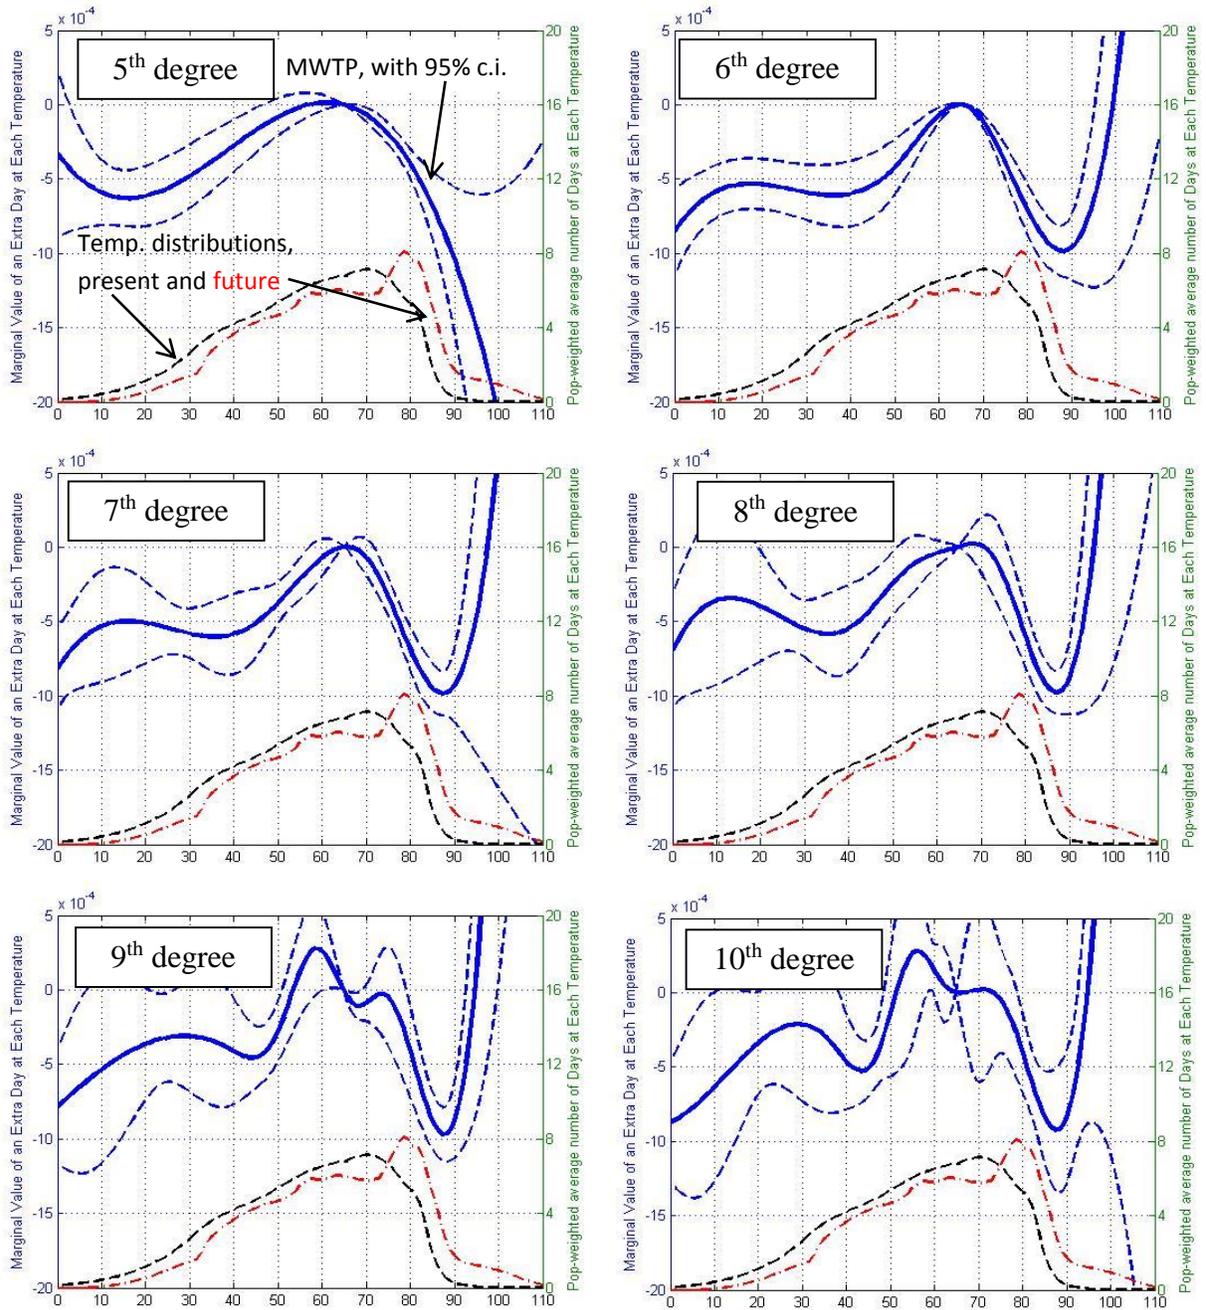


Figure 3: Present and projected future January and July mean temperatures



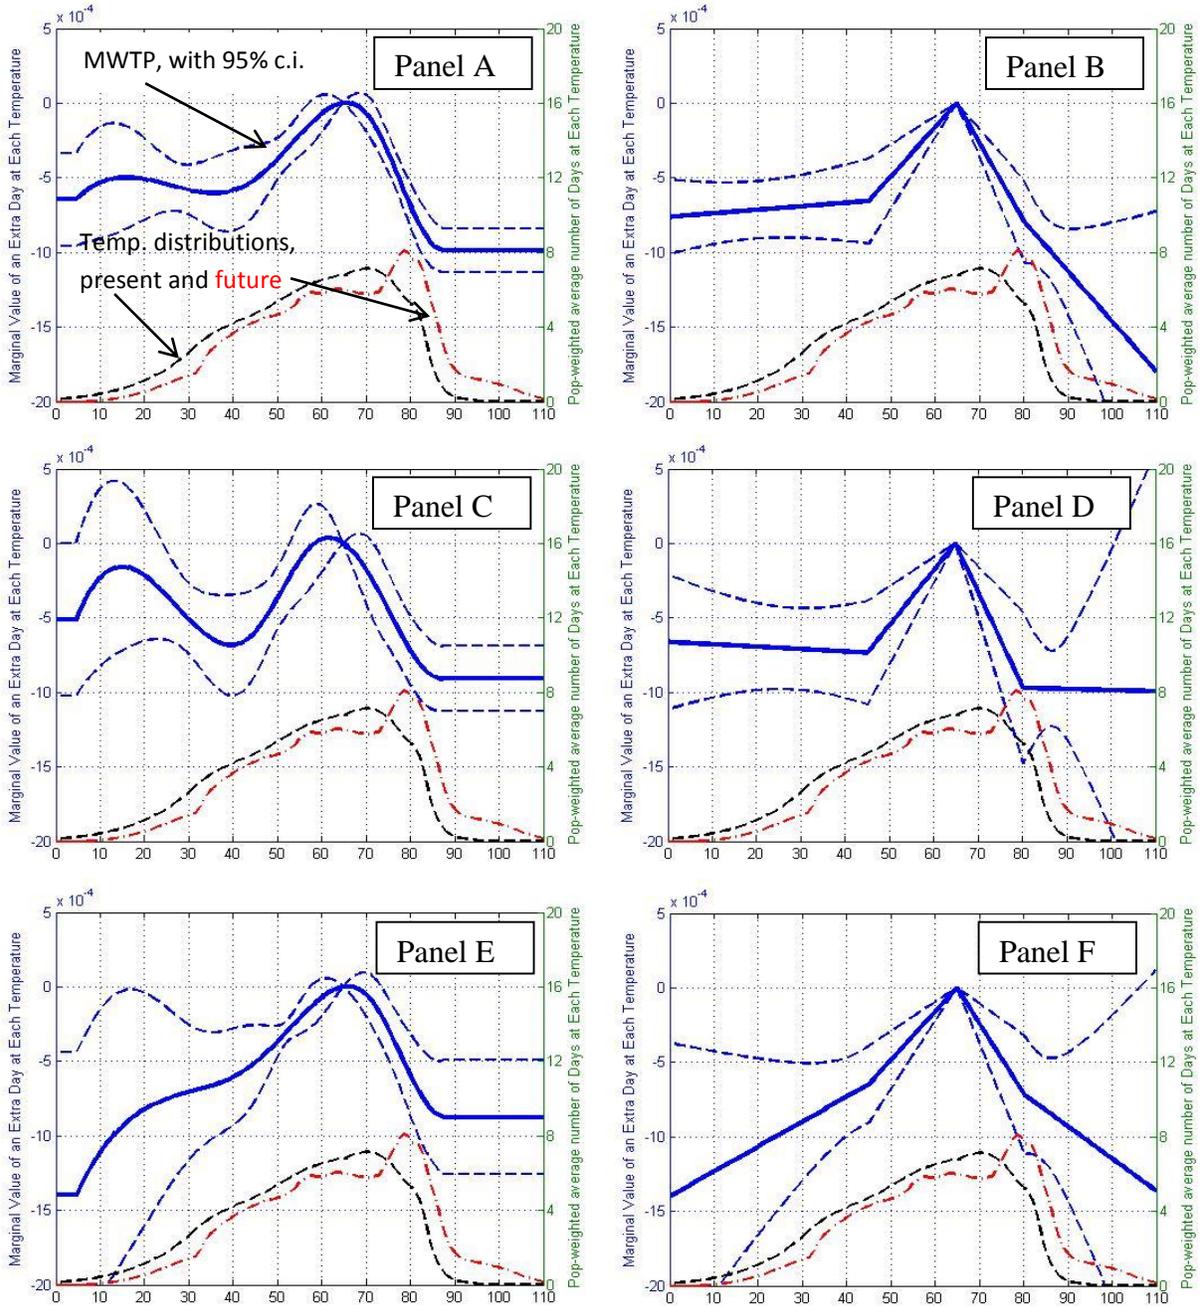
Notes: Present-day temperatures are 1970-1999 average. Projections use the A2 scenario for 2070-2099.

Figure 4: Estimated WTP for temperature, homogenous preference cubic spline models



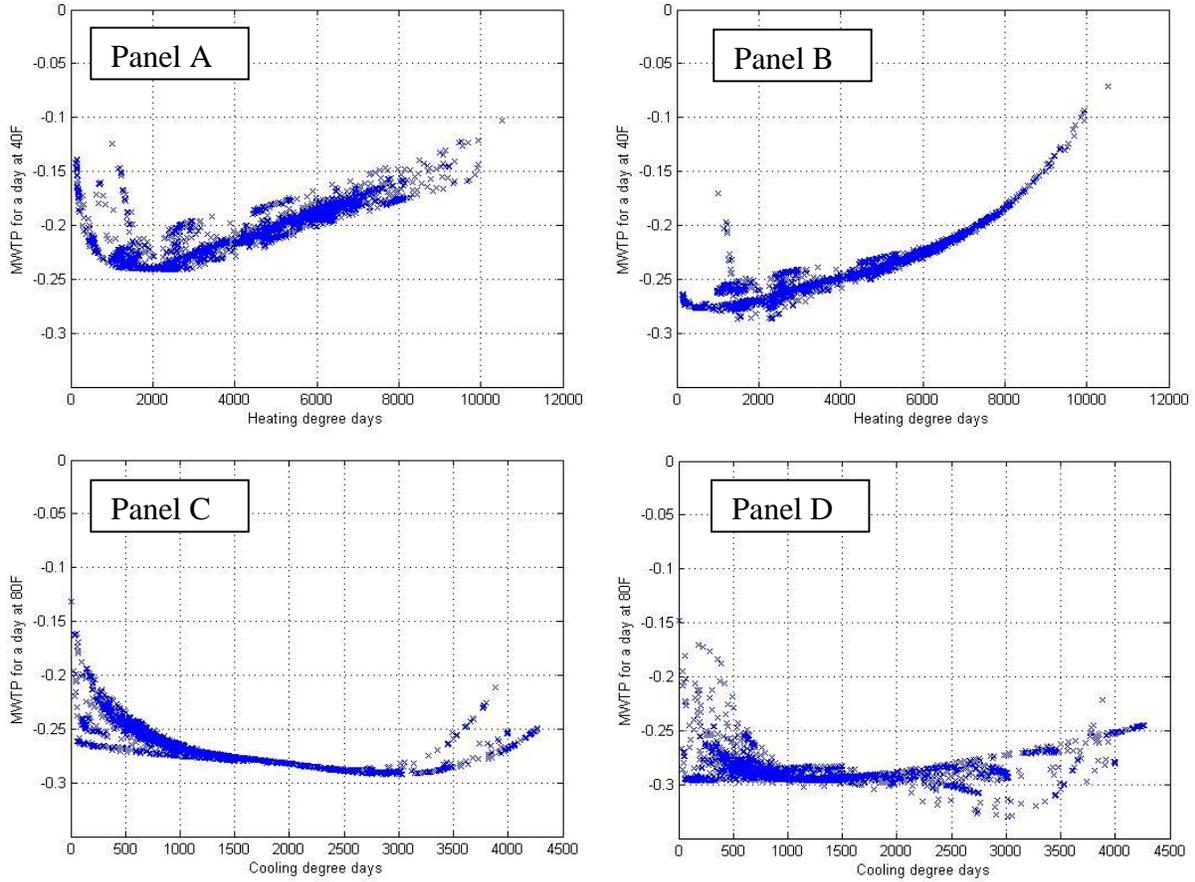
Notes: All specifications use the full set of geographic, “other weather”, and demographic controls. MWTPs are expressed as a fraction of income and normalized to zero at 65°F. Each observation (PUMA) is weighted by its population in the regressions. Standard errors are clustered on state. “Present” temperature distribution denotes 1970-1999, and “future” denotes 2070-2099. Both distributions are population-weighted U.S. averages.

Figure 5: Estimated WTP for temperature, homogenous preference models



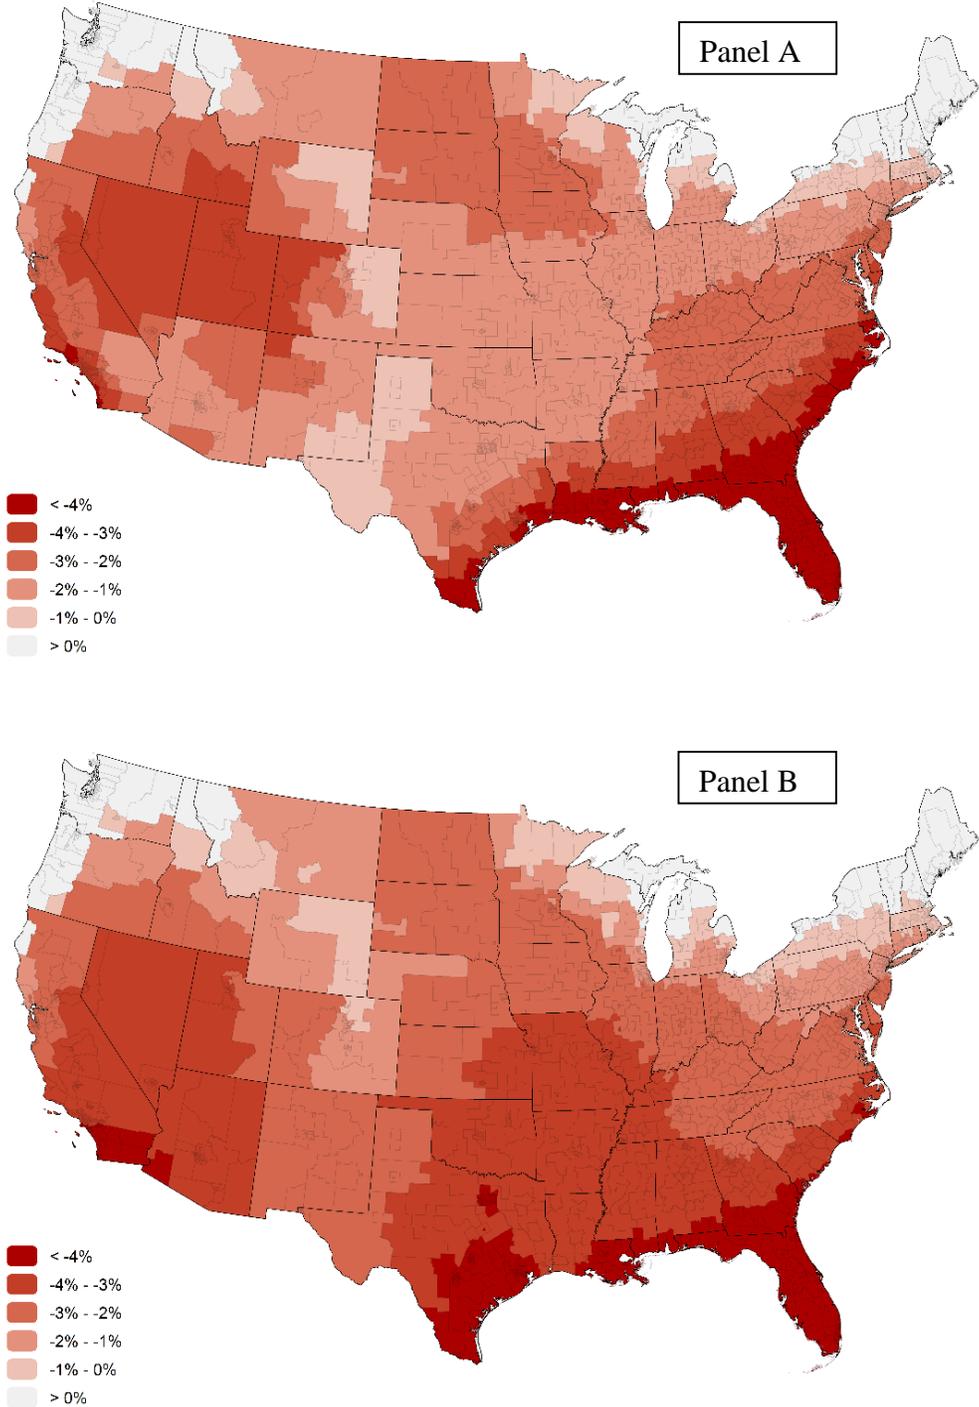
Notes: Panels A, C, and E, use a 7th degree cubic spline WTP model, restricted to be constant over the extreme 1% of the temperature distribution. Panels B, D, and E use a four-piece linear spline. Panels A and B include all geographic, “other weather”, and demographic controls. Panels C and D omit demographic controls. Panels E and F include all controls and state fixed effects. MWTPs are expressed as a fraction of income and normalized to zero at 65°F. Each observation (PUMA) is weighted by its population in the regressions. Standard errors are clustered on state. “Present” temperature distribution denotes 1970-1999, and “future” denotes 2070-2099. Both distributions are population-weighted U.S. averages.

Figure 6: Estimated distribution of MWTPs for heat and cold from the heterogeneous preferences models



Notes: MWTPs are expressed as a fraction of income, scaled up by 365, and normalized to zero at 65°F. Each plotted point denotes a PUMA-specific estimate of MWTP. Panels A and C use the restricted 7th degree cubic spline model, and panels B and D use the four-piece linear spline. Models used in all panels include the geographic, “other weather”, and demographic controls. Each observation (PUMA) is weighted by its population in the regressions.

Figure 7: Predicted changes in QOL as percent of income under the A2 scenario for 2070-2099, with the homogenous preference model WTP estimates



Notes: Panel A uses the restricted 7th degree cubic spline WTP model, per panel A of figure 5 and specification I of tables 2 and 4. Panel B uses the four-piece linear spline WTP model, per panel B of figure 5 and specification II of tables 2 and 4.

Table 1: Descriptive statistics for primary dataset.
Sample consists of 2057 Public Use Microdata Areas (PUMAs)

	Mean	Std. Dev	10th percentile	90th percentile
<i>Temperature data, 1970-1999 average</i>				
Average annual heating degree days (1000s)	4.384	2.204	1.326	7.009
Average annual cooling degree days (1000s)	1.290	0.929	0.374	2.762
<i>Temperature data, 2070-2099 projected (CCSM A2)</i>				
Projected avg. annual heating degree days (1000s)	2.974	1.729	0.527	5.024
Projected avg. annual cooling degree days (1000s)	2.547	1.190	1.304	4.413
<i>Other climate data, 1970-1999 average</i>				
Average annual precipitation (inches)	39.26	14.09	16.25	53.85
Average annual relative humidity (%)	63.58	8.10	53.31	70.52
Average annual sunshine (% of available daylight)	60.18	8.64	49.74	73.08
<i>Other climate data, 2070-2099 projected (CCSM A2)</i>				
Projected avg. annual precipitation (inches)	41.54	15.23	15.28	56.40
Projected avg. annual relative humidity (%)	62.66	8.74	51.18	69.63
Projected avg. annual sunshine (% of available daylight)	61.37	8.56	51.31	72.85
<i>Geographic data</i>				
Area in square miles	1,436	4,255	17	3,355
Distance from centroid of puma to ocean (miles)	250.1	272.3	4.3	729.2
Distance from centroid of puma to Great Lake (miles)	763.2	715.4	54.0	2128.4
Average land slope (degrees)	1.677	2.131	0.191	4.270
<i>Demographic data (2000 census)</i>				
Weighted population density (people per sq. mile)	5,466	11,997	360	9,981
Percent high school graduates	83.9%	8.8%	72.1%	93.3%
Percent of population with bachelors degree	24.1%	12.4%	11.3%	41.0%
Percent of population with graduate degree	8.7%	5.6%	3.7%	20.3%
Average age	35.8	3.1	32.0	39.3
Percent hispanic	8.9%	13.5%	0.6%	25.8%
Percent black	12.5%	17.5%	0.7%	36.5%
Population (1000s)	135.9	32.9	103.8	182.2
Quality of life differential (in logs)	-0.001	0.784	-0.089	0.101

Table 2: Estimation results for the linear, homogenous preferences model.
Dependent variable is quality of life (QOL) differential as fraction of income

	I	II	III	IV	V	VI
Temperature function	Restricted cubic spline	Linear spline	Restricted cubic spline	Linear spline	Restricted cubic spline	Linear spline
MWTP for a day at 40F (x365)	-0.215 (0.050)	-0.244 (0.047)	-0.250 (0.063)	-0.265 (0.058)	-0.223 (0.061)	-0.270 (0.049)
MWTP for a day at 80F (x365)	-0.254 (0.028)	-0.291 (0.051)	-0.265 (0.043)	-0.354 (0.095)	-0.220 (0.051)	-0.263 (0.074)
Average annual precipitation (inches, x10)	0.011 (0.002)	0.011 (0.002)	0.005 (0.003)	0.005 (0.003)	0.011 (0.003)	0.011 (0.003)
Average annual relative humidity (fraction)	-0.040 (0.047)	-0.067 (0.038)	-0.041 (0.081)	-0.050 (0.063)	-0.049 (0.060)	-0.060 (0.065)
Average annual sunshine (fraction of all daytime hours)	0.151 (0.051)	0.147 (0.049)	0.120 (0.080)	0.115 (0.068)	0.062 (0.092)	0.080 (0.093)
Average land slope (degrees)	0.004 (0.001)	0.004 (0.001)	0.006 (0.002)	0.006 (0.002)	0.003 (0.001)	0.003 (0.001)
Inverse distance to ocean (miles)	0.238 (0.034)	0.235 (0.034)	0.397 (0.043)	0.398 (0.044)	0.213 (0.040)	0.204 (0.038)
Inverse distance to Great Lakes (miles)	0.174 (0.126)	0.198 (0.122)	0.078 (0.220)	0.128 (0.207)	0.091 (0.132)	0.106 (0.130)
Squared inverse distance to ocean (miles)	-0.200 (0.027)	-0.199 (0.027)	-0.314 (0.065)	-0.314 (0.063)	-0.192 (0.033)	-0.185 (0.033)
Squared inverse distance to Great Lakes (miles)	-0.158 (0.140)	-0.186 (0.134)	0.051 (0.240)	-0.006 (0.221)	-0.088 (0.146)	-0.103 (0.144)
Log of weighted population density	0.006 (0.003)	0.007 (0.003)	-	-	0.003 (0.002)	0.004 (0.002)
Fraction high school graduates	0.156 (0.048)	0.154 (0.046)	-	-	0.134 (0.045)	0.131 (0.043)
Fraction bachelors degrees	0.314 (0.059)	0.298 (0.063)	-	-	0.338 (0.058)	0.341 (0.059)
Fraction graduate degrees	-0.011 (0.095)	0.019 (0.105)	-	-	-0.027 (0.091)	-0.032 (0.095)
Average age (x100)	0.204 (0.040)	0.204 (0.039)	-	-	0.220 (0.033)	0.209 (0.036)
Fraction hispanic	0.089 (0.033)	0.082 (0.035)	-	-	0.092 (0.027)	0.089 (0.028)
Fraction black	-0.064 (0.008)	-0.068 (0.008)	-	-	-0.057 (0.009)	-0.060 (0.008)
State fixed effects	N	N	N	N	Y	Y
Number of observations	2057	2057	2057	2057	2057	2057
R ²	0.798	0.796	0.386	0.382	0.830	0.830

Notes: The unit of observation is a PUMA. All regressions are weighted by population in each PUMA. The linear splines are four-piece with breakpoints at 45F, 65F, and 80F, and the restricted cubic spline is 7th degree with a constant MWTP restriction for the lower and upper 1% of the population-weighted U.S. average temperature distribution. MWTPs shown for a day at 40F or 80F are relative to 65F and expressed as a fraction of income. Parenthetical values indicate standard errors clustered on state.

Table 3: Estimation results for the heterogeneous preferences model.
Dependent variable is quality of life (QOL) differential as fraction of income

	I	II	III	IV	V	VI
Temperature function	Restricted cubic spline	Linear spline	Restricted cubic spline	Linear spline	Restricted cubic spline	Linear spline
Average MWTP for a day at 40F (x365)	-0.203 (0.046)	-0.239 (0.044)	-0.240 (0.059)	-0.274 (0.054)	-0.221 (0.051)	-0.268 (0.039)
Average MWTP for a day at 80F (x365)	-0.265 (0.025)	-0.287 (0.047)	-0.337 (0.046)	-0.416 (0.071)	-0.236 (0.043)	-0.269 (0.059)
Geog & other weather controls	Y	Y	Y	Y	Y	Y
Demographic controls	Y	Y	N	N	Y	Y
State fixed effects	N	N	N	N	Y	Y

Notes: The unit of observation is a PUMA. All regressions are weighted by population in each PUMA. The linear splines are four-piece with breakpoints at 45F, 65F, and 80F, and the restricted cubic spline is 7th degree with a constant MWTP restriction for the lower and upper 1% of the population-weighted U.S. average temperature distribution. MWTPs shown for a day at 40F or 80F are relative to 65F, expressed as a fraction of income, and are income-weighted averages across all PUMAs. Parenthetical values indicate standard errors wild cluster bootstrapped on state using 200 draws.

Table 4: Projected U.S. average amenity value impacts from climate change for 2070-2099

	I	II	III	IV	V	VI
	Homogenous preferences			Heterogeneous preferences		
Temperature function	Restricted cubic spline	Linear spline	Restricted cubic spline	Restricted cubic spline	Linear spline	Restricted cubic spline
A2 scenario change in welfare as percent of income	-2.28% (0.25%)	-2.50% (0.62%)	-0.97% (0.79%)	-2.79% (0.27%)	-3.91% (0.61%)	-1.02% (0.83%)
Component driven by warmer winters	2.12% (0.26%)	2.70% (0.38%)	2.94% (0.80%)	2.12% (0.24%)	2.77% (0.35%)	3.05% (0.80%)
Component driven by hotter summers	-4.85% (0.36%)	-5.68% (0.66%)	-4.27% (0.89%)	-5.36% (0.31%)	-7.16% (0.54%)	-4.44% (0.89%)
Component driven by other weather changes	0.45% (0.09%)	0.48% (0.08%)	0.37% (0.14%)	0.45% (0.09%)	0.48% (0.08%)	0.37% (0.14%)
Percent of population losing	93.6% (2.6%)	94.2% (2.2%)	70.7% (17.7%)	95.7% (2.4%)	95.1% (2.0%)	70.3% (17.4%)
AIFI change in welfare as percent of income	-2.91% (0.31%)	-3.23% (0.82%)	-1.47% (0.95%)	-3.53% (0.34%)	-5.10% (0.83%)	-1.53% (1.00%)
Geog & other weather controls	Y	Y	Y	Y	Y	Y
Demographic controls	Y	Y	Y	Y	Y	Y
State fixed effects	N	N	Y	N	N	Y

Notes: Parenthetical values indicate standard errors wild cluster bootstrapped on state using 200 draws.

Table 5: Preference estimates and projected U.S. average amenity value impacts from climate change for 2070-2099. Heterogeneous preferences model with alternative quality of life (QOL) measures

	I	II	III	IV	V	VI	VII	VIII
QOL measure	Reference case, 7th degree cubic spline	Wages by place of residence	Wages only (negative)	Prices only	QOL = 0.25p - 1.00w	Non-college graduates only	College graduates only	Migration / selection adjustment
<u>Demographics, no state FE</u>								
Average MWTP for a day at 40F (x365)	-0.203 (0.046)	-0.125 (0.040)	0.193 (0.115)	-0.726 (0.180)	0.030 (0.084)	-0.227 (0.043)	-0.149 (0.056)	-0.208 (0.050)
Average MWTP for a day at 80F (x365)	-0.265 (0.025)	-0.177 (0.036)	0.573 (0.118)	-1.455 (0.159)	0.243 (0.115)	-0.237 (0.023)	-0.293 (0.032)	-0.300 (0.034)
A2 change in welfare as percent of income	-2.79% (0.27%)	-2.30% (0.30%)	5.20% (1.32%)	-14.86% (1.69%)	1.47% (0.92%)	-2.27% (0.28%)	-3.47% (0.31%)	-3.01% (0.29%)
<u>Demographics, state FE</u>								
Average MWTP for a day at 40F (x365)	-0.221 (0.051)	-0.178 (0.045)	-0.147 (0.057)	-0.403 (0.196)	-0.231 (0.047)	-0.242 (0.048)	-0.172 (0.063)	-0.221 (0.055)
Average MWTP for a day at 80F (x365)	-0.236 (0.043)	-0.210 (0.047)	0.081 (0.084)	-0.747 (0.236)	-0.115 (0.065)	-0.231 (0.038)	-0.207 (0.060)	-0.260 (0.055)
A2 change in welfare as percent of income	-1.02% (0.83%)	-1.82% (0.58%)	-0.57% (1.24%)	-2.95% (3.18%)	-2.46% (0.88%)	-0.96% (0.77%)	-0.53% (1.00%)	-0.94% (0.89%)

Notes : The unit of observation is a PUMA. All regressions are weighted by population in each PUMA. All specifications use a 7th degree restricted cubic spline with a constant MWTP restriction for the lower and upper 1% of the population-weighted U.S. average temperature distribution. All regressions include geographic, other weather, and demographic controls. MWTPs shown for a day at 40F or 80F are relative to 65F, expressed as a fraction of income, and are income-weighted averages across all PUMAs. Welfare impacts shown assume that households' utility functions are linear in characteristics. Column V uses wages by residence. Column VIII adjusts QOL for selective migration per Dahl (2002), as discussed in sections IB and VI. Parenthetical values indicate standard errors wild cluster bootstrapped on state using 200 draws.

NOT FOR PUBLICATION

**ONLINE APPENDICES FOR “CLIMATE AMENITIES, CLIMATE CHANGE,
AND AMERICAN QUALITY OF LIFE”**

by

David Albouy, Walter Graff, Ryan Kellogg, and Hendrik Wolff

Appendix 1: Supplementary figures and tables

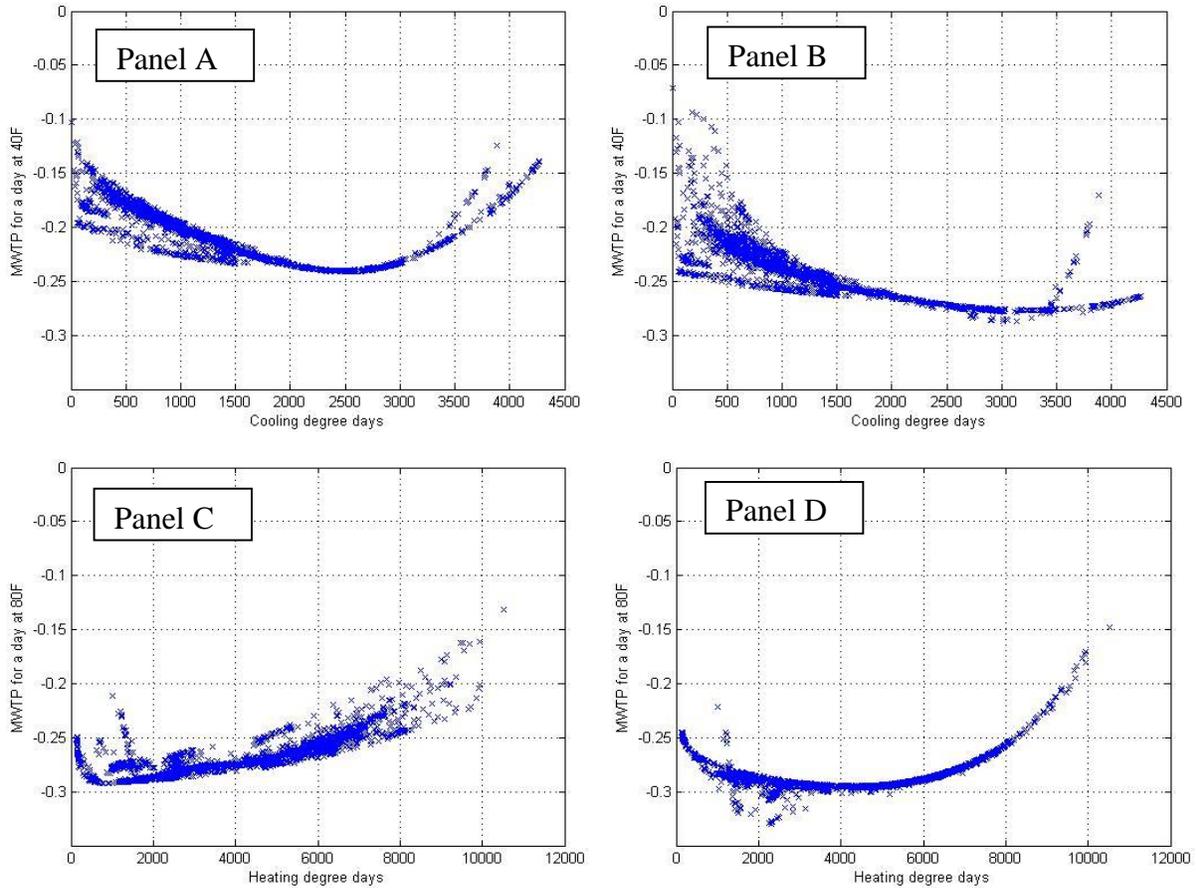
Appendix 2: Data construction

Appendix 3: Utility function concavity

Appendix 4: Set identification

Appendix 1: Supplementary Figures and Tables

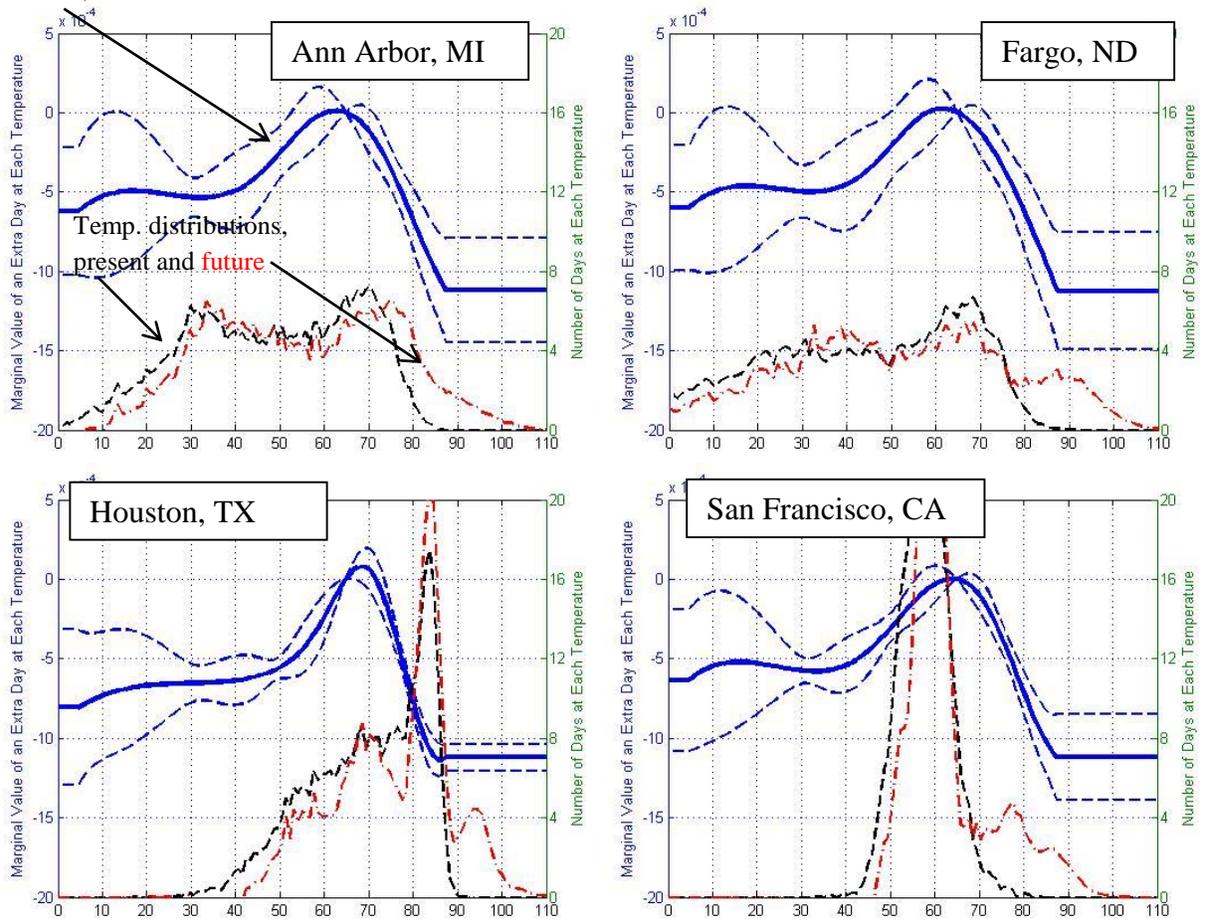
Figure A1.1: Estimated MWTPs for heat and cold from the heterogeneous preferences models; MWTP to avoid cold (40°F) plotted against CDD and MWTP to avoid heat (80°F) plotted against HDD



Notes: MWTPs are expressed as a fraction of income, scaled up by 365, and normalized to zero at 65°F. Each plotted point denotes a PUMA-specific estimate of MWTP. Panels A and C use the restricted 7th degree cubic spline model, and panels B and D use the four-piece linear spline. Models used in all panels include the geographic, “other weather”, and demographic controls. Each observation (PUMA) is weighted by its population in the regressions.

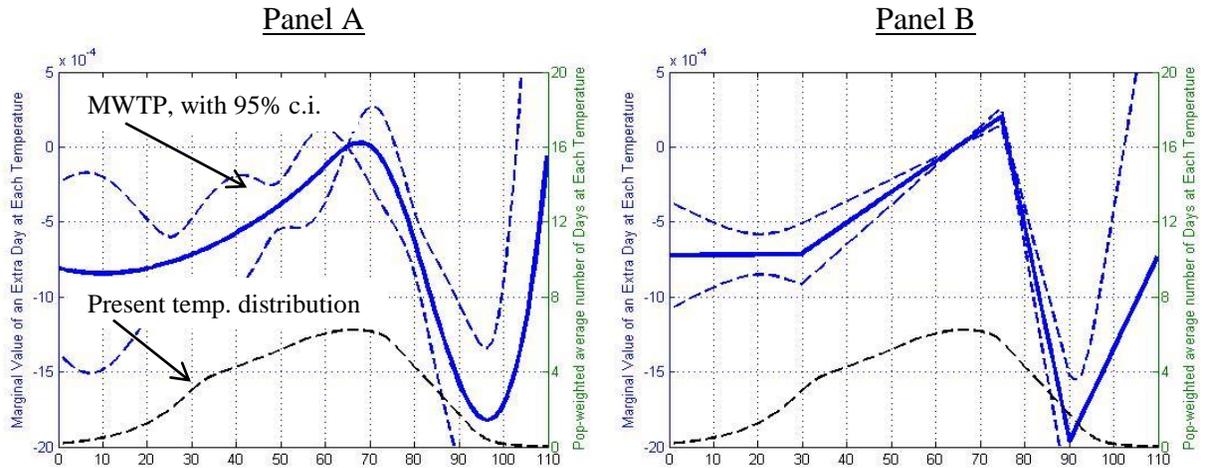
Figure A1.2: Estimated WTP for temperature and four specific locations, heterogeneous preference model

MWTP, with 95% c.i.



Notes: All panels use a 7th degree cubic spline WTP model, restricted to be constant over the extreme 1% of the temperature distribution, and include all geographic, “other weather”, and demographic controls. MWTPs are expressed as a fraction of income and normalized to zero at 65°F. Each observation (PUMA) is weighted by its population in the regressions. Standard errors are clustered on state. “Present” temperature distribution denotes 1970-1999, and “future” denotes 2070-2099.

Figure A1.3: Estimated WTP for temperature, homogenous preference model, local temperature distributions include within-day variation



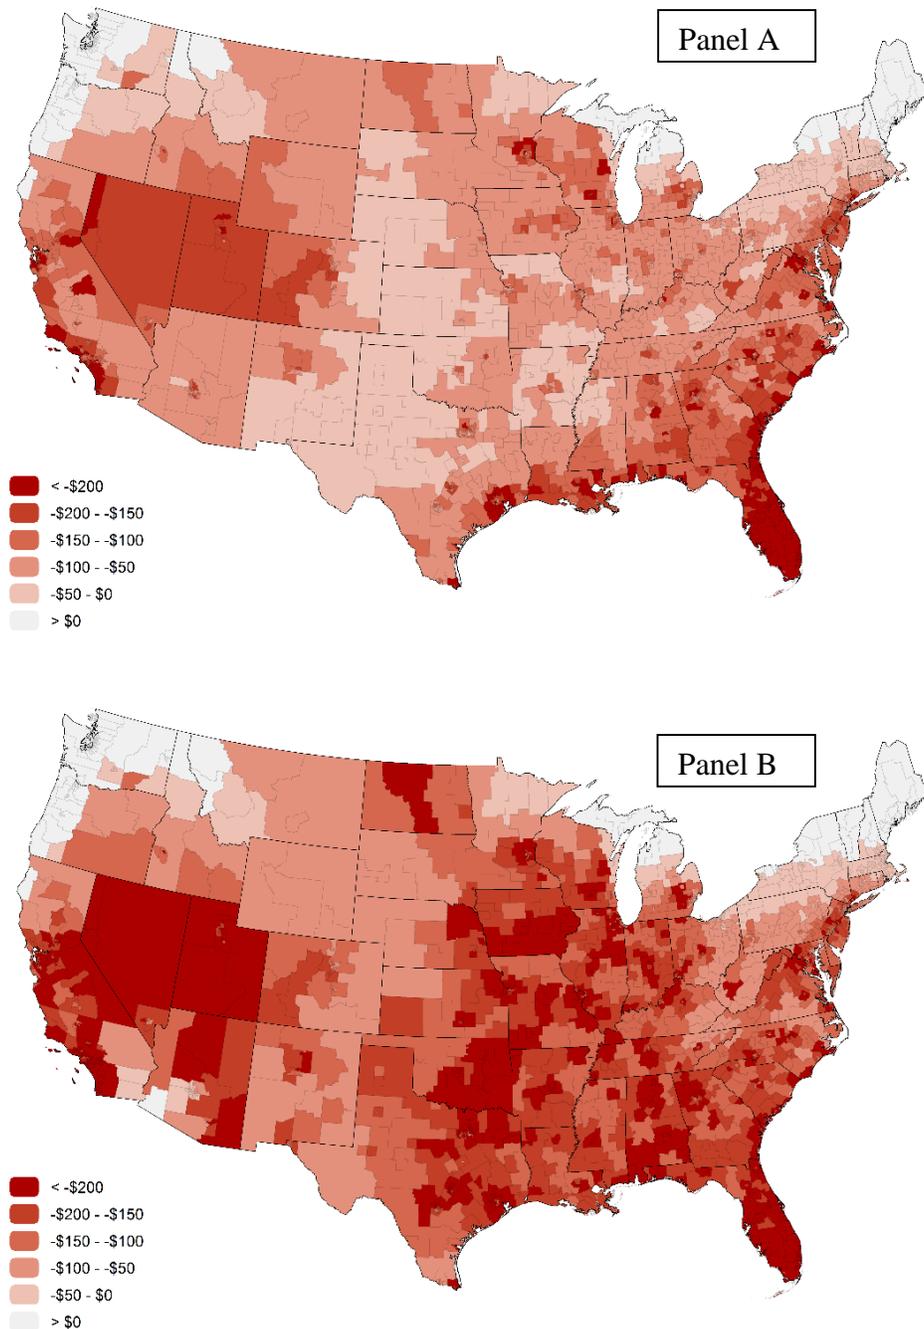
Notes: The “present” temperature distribution in both panels incorporates intraday variation between the daily high and low temperatures at each PRISM gridpoint. Rather than bin temperatures based on the daily mean, we fit a sinusoidal curve to each gridpoint-day’s high and low temperatures and calculate the fraction of the day in which the temperature is in each bin. We then aggregate across days and across gridpoints within each Puma as described in the main text and appendix 2.

Panel A uses a 7th degree cubic spline WTP model, and panel B uses a four-piece linear spline. The cubic spline knots are spaced evenly on the cdf of the population-weighted average temperature distribution, including intraday variation. The linear spline knots are at 20°F, 75°F, and 90°F.

Both panels include all geographic, “other weather”, and demographic controls. MWTPs are expressed as a fraction of income and normalized to zero at 65°F. Each observation (PUMA) is weighted by its population in the regressions. Standard errors are clustered on state.

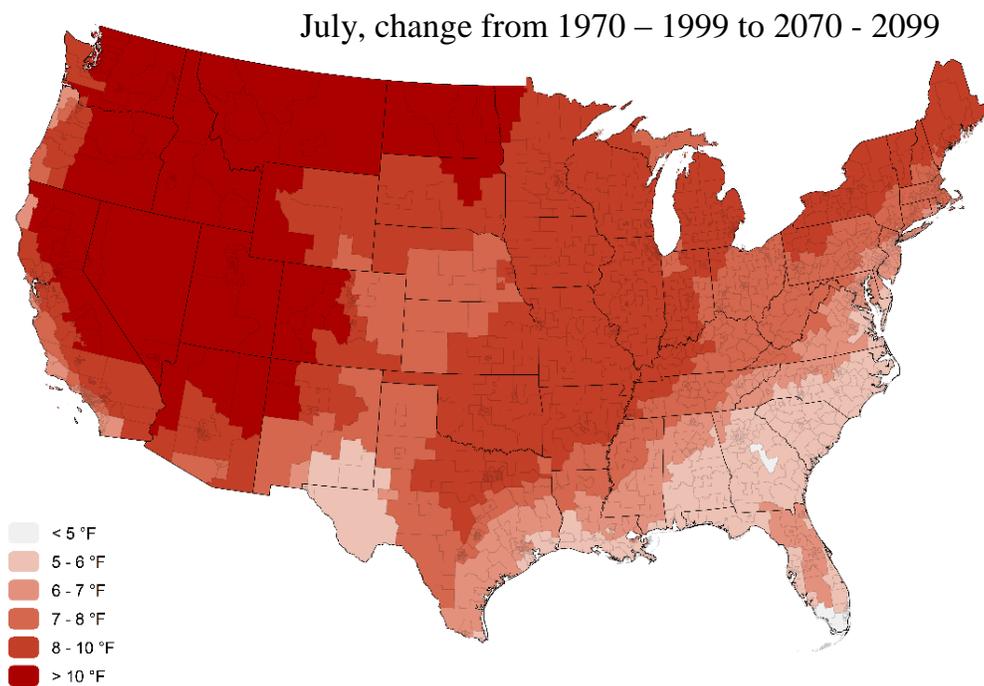
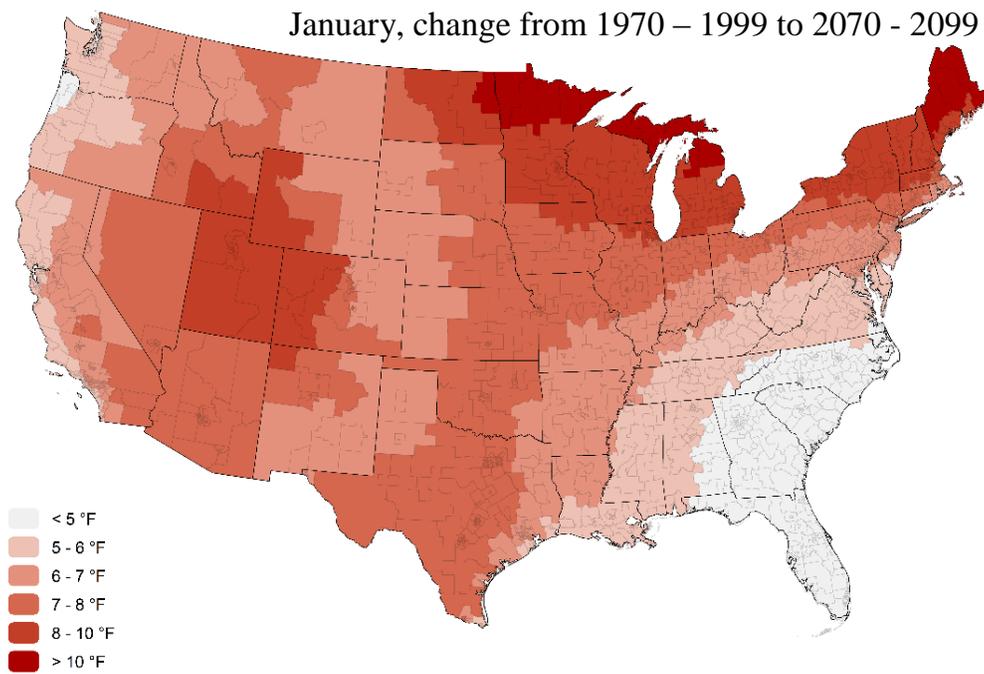
In panel B, the null hypothesis that the magnitudes of the slopes on either side of the 75°F knot are equal is rejected with $p < 0.001$. The null hypotheses that there is no kink at either 20°F or 90°F are rejected with $p = 0.063$ and 0.006 , respectively. We fail to reject the null hypotheses that WTP is constant to the left of 20°F and to the right of 90°F with $p = 0.972$ and $p = 0.317$, respectively.

Figure A1.4: Predicted changes in QOL in millions of 2000 U.S. dollars, under the A2 scenario for 2070-2099, with the heterogeneous preference model MWTP estimates



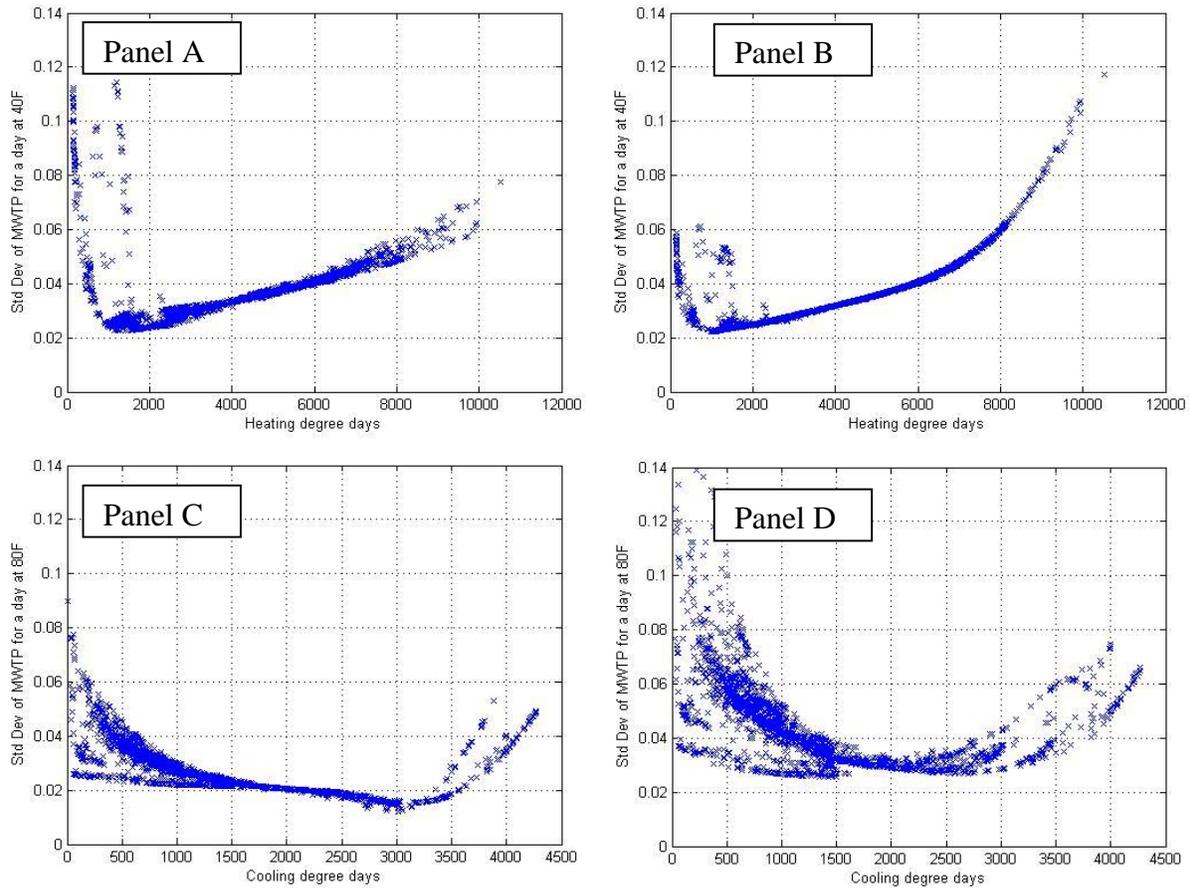
Notes: Panel A uses the restricted 7th degree cubic spline WTP model, per panels A and C of figure 6, specification I of table 3, and specification IV of table 4. Panel B uses the four-piece linear spline WTP model, per panels B and D of figure 6, specification II of table 3, and specification V of table 4. Total 2000 personal income for the contiguous 48 states was \$8.576 trillion. Total losses across all PUMAs are \$239 billion in panel A and \$336 billion in panel B.

Figure A1.5: Projected changes to future January and July mean temperatures



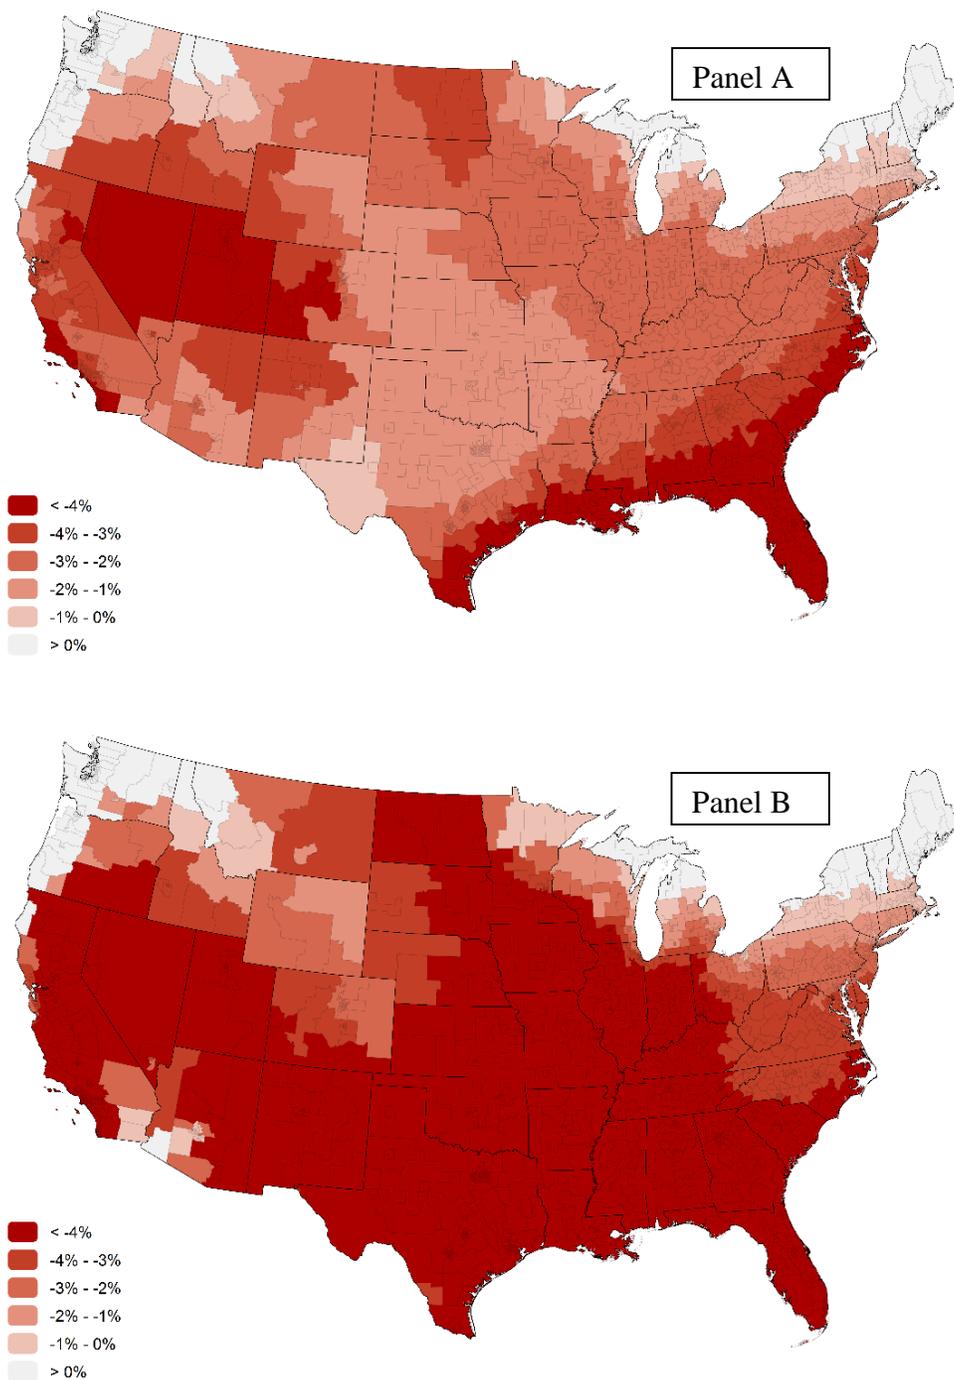
Notes: Projections use the A2 scenario for 2070-2099.

Figure A1.6: Estimated distribution of standard deviation of MWTPs for heat and cold from the heterogeneous preferences models



Notes: MWTPs are expressed as a fraction of income, scaled up by 365, and normalized to zero at 65°F. Each plotted point denotes the standard deviation of a PUMA-specific estimate of MWTP. Panels A and C use the restricted 7th degree cubic spline model, and panels B and D use the four-piece linear spline. Models used in all panels include the geographic, “other weather”, and demographic controls. Each observation (PUMA) is weighted by its population in the regressions.

Figure A1.7: Predicted changes in QOL as percent of income under the A2 scenario for 2070-2099, with the heterogeneous preference model MWTP estimates



Notes: Panel A uses the restricted 7th degree cubic spline WTP model, per panels A and C of figure 6, specification I of table 3, and specification IV of table 4. Panel B uses the four-piece linear spline WTP model, per panels B and D of figure 6, specification II of table 3, and specification V of table 4.

Table A1.1, part 1: Alternative specifications for "geographic" and "other weather" controls

	Reference Case	R1	R2	R3	R4	R5	R6	R7	R8
	I	R1	R2	R3	R4	R5	R6	R7	R8
	Restricted cubic spline as in Specification I of tables 2-4	As in I, plus a cubic added in inverse distance to water bodies	As in I, plus a quadratic added in the average PUMA mountainousness	As in I, except "Relative Humidity" replaced by "Dew Point"	As in I, except "Relative Humidity" replaced by "Relative Humidity in summer months only"	As in I, except "Relative Humidity" replaced by "Dew Point in summer months only"	As in I, except all "other weather" variables split up into summer and winter months variables*	As in R6, except "Relative Humidity" replaced by "Dew Point"	As in I, except "Sunshine" calculated as interpolated using all available stations (and not just the four closest stations)
<i>Panel a: Preference estimates and projected U.S. average amenity value impacts from climate change for 2090-2100 using the homogeneous preference models</i>									
MWTP for a day at 40F (x365)	-0.215 (0.050)	-0.196 (0.052)	-0.221 (0.048)	-0.219 (0.069)	-0.251 (0.060)	-0.260 (0.057)	-0.205 (0.062)	-0.078 (0.079)	-0.208 (0.052)
MWTP for a day at 80F (x365)	-0.254 (0.028)	-0.230 (0.027)	-0.259 (0.029)	-0.248 (0.030)	-0.276 (0.032)	-0.236 (0.020)	-0.238 (0.030)	-0.225 (0.020)	-0.251 (0.028)
A2 scenario change in welfare as percent of income	-2.28% (0.25%)	-2.02% (0.23%)	-2.34% (0.27%)	-2.31% (0.30%)	-2.61% (0.35%)	-3.24% (0.39%)	-2.35% (0.27%)	-2.57% (0.40%)	-2.27% (0.24%)
A1FI scenario change in welfare as percent of income	-2.91% (0.31%)	-2.58% (0.28%)	-2.99% (0.33%)	-2.95% (0.38%)	-3.34% (0.42%)	-4.10% (0.47%)	-3.02% (0.36%)	-3.23% (0.51%)	-2.90% (0.30%)

Panel b: Preference estimates and projected U.S. average amenity value impacts from climate change for 2090-2100 using the heterogeneous preference models

Average MWTP for a day at 40F (x365)	-0.203 (0.046)	-0.185 (0.048)	-0.208 (0.044)	-0.207 (0.067)	-0.239 (0.060)	-0.253 (0.055)	-0.196 (0.061)	-0.076 (0.075)	-0.196 (0.048)
Average MWTP for a day at 80F (x365)	-0.265 (0.025)	-0.236 (0.026)	-0.272 (0.026)	-0.258 (0.031)	-0.287 (0.031)	-0.236 (0.021)	-0.247 (0.032)	-0.213 (0.024)	-0.261 (0.025)
A2 scenario change in welfare as percent of income	-2.79% (0.27%)	-2.49% (0.25%)	-2.87% (0.28%)	-2.81% (0.34%)	-3.09% (0.36%)	-3.50% (0.44%)	-2.75% (0.35%)	-2.68% (0.45%)	-2.79% (0.27%)
A1FI scenario change in welfare as percent of income	-3.53% (0.34%)	-3.15% (0.31%)	-3.63% (0.36%)	-3.56% (0.43%)	-3.91% (0.45%)	-4.41% (0.53%)	-3.52% (0.44%)	-3.38% (0.58%)	-3.52% (0.34%)

Notes: The unit of observation is a PUMA, except for R13, where the unit of observation is a MSA. MWTPs shown for a day at 40F or 80F are relative to 65F, expressed as a fraction of income, and are income-weighted averages across all observations. Except for R11, all regressions are weighted by population. Parenthetical values indicate standard errors cluster bootstrapped on state using 200 draws (except for the first two rows, which are not bootstrapped). Welfare impacts shown assume that households' utility functions are linear in characteristics.

Dew point is highly collinear with temperature, so that in specifications R3, R5, and R7 the MWTPs for temperature decrease as this MWTP is loaded onto the dew point coefficient instead. The welfare predictions do not change substantially because the forecasted increase in dew point is similar to the forecasted increase in temperature.

* This split effectively adds three variables to the specification

Table A1.1, part 2: Alternative specifications and standard errors

Reference Case	I					R9					R10					R11					R12					R13					R14					R15				
	Restricted cubic spline as in Specification I of tables 2-4					As in I, plus squares of all demographic variables added					As in I plus "Census Division Fixed Effects" added					As in I, except regressions are unweighted					As in I, except temperature modeled by each 10 degree bin***					As in I, except observations based on MSA level					As in I, except standard errors clustered by Census Division					As in I, except 'Superstar' MSAs dropped from the dataset***				
MWTP for a day at 40F (x365)	-0.215	(0.050)	-0.215	(0.039)	-0.208	(0.053)	-0.214	(0.052)	-0.226	(0.083)	-0.202	(0.052)	-0.215	(0.033)	-0.138	(0.055)																								
MWTP for a day at 80F (x365)	-0.254	(0.028)	-0.243	(0.033)	-0.209	(0.028)	-0.255	(0.031)	-0.267	(0.038)	-0.213	(0.036)	-0.254	(0.029)	-0.182	(0.033)																								
A2 scenario change in welfare as percent of income	-2.28%	(0.25%)	-1.94%	(0.30%)	-1.84%	(0.35%)	-2.28%	(0.27%)	-2.29%	(0.53%)	-1.88%	(0.36%)	-2.28%	(0.32%)	-1.68%	(0.47%)																								
AIFF scenario change in welfare as percent of income	-2.91%	(0.31%)	-2.53%	(0.36%)	-2.46%	(0.43%)	-2.92%	(0.32%)	-2.90%	(0.71%)	-2.41%	(0.45%)	-2.91%	(0.39%)	-2.16%	(0.58%)																								

Panel a: Preference estimates and projected U.S. average amenity value impacts from climate change for 2070-2099 using the homogenous preference models

Panel b: Preference estimates and projected U.S. average amenity value impacts from climate change for 2070-2099 using the heterogeneous preference models

Average MWTP for a day at 40F (x365)	-0.203	(0.046)	-0.203	(0.036)	-0.203	(0.049)	-0.202	(0.047)	-0.175	(0.100)	-0.192	(0.045)	-0.203	(0.032)	-0.138	(0.041)
Average MWTP for a day at 80F (x365)	-0.265	(0.025)	-0.255	(0.033)	-0.210	(0.024)	-0.270	(0.028)	-0.252	(0.069)	-0.202	(0.033)	-0.265	(0.023)	-0.190	(0.027)
A2 scenario change in welfare as percent of income	-2.79%	(0.27%)	-2.53%	(0.29%)	-1.99%	(0.47%)	-2.85%	(0.29%)	-3.23%	(0.46%)	-2.46%	(0.40%)	-2.79%	(0.35%)	-2.28%	(0.44%)
AIFF scenario change in welfare as percent of income	-3.53%	(0.34%)	-3.27%	(0.36%)	-2.63%	(0.60%)	-3.62%	(0.36%)	-4.07%	(0.62%)	-3.09%	(0.52%)	-3.53%	(0.44%)	-2.86%	(0.55%)

Notes: The unit of observation is a PUMA, except for R13, where the unit of observation is a MSA. MWTPs shown for a day at 40F or 80F are relative to 65F, expressed as a fraction of income, and are income-weighted averages across all observations. Except for R11, all regressions are weighted by population. Parenthetical values indicate standard errors cluster bootstrapped on state using 200 draws (except for the first two rows, which are not bootstrapped). Welfare impacts shown assume that households' utility functions are linear in characteristics.

** The temperature bins are defined in the ranges (-inf, 35], (35,45], (45,55], ..., (75,85], and (85,+inf).

*** Superstar msas are those listed in Gyorko, Mayer, and Sinai (2013), table A3, for the year 2000. These msas account for 442 of the 2,057 PUMAs in the overall sample.

Table A1.2: Correlations between HDD, CDD, and control variables

	HDD	CDD
<i>Geographic data</i>		
Distance from centroid of puma to ocean (miles)	0.568	-0.291
Distance from centroid of puma to Great Lake (miles)	-0.521	0.156
Average land slope (degrees)	0.199	-0.360
<i>Demographic data (2000 census)</i>		
Weighted population density (people per sq. mile)	-0.031	-0.049
Percent high school graduates	0.305	-0.247
Percent of population with bachelors degree	0.068	-0.141
Percent of population with graduate degree	0.073	-0.140
Average age	0.120	-0.036
Percent hispanic	-0.420	0.295
Percent black	-0.203	0.212

Appendix 2: Data Construction

1 Overview

Here we document the construction of the dataset used to estimate the models in the paper. Specifically, we collect the following data for each PUMA:

- historical temperature
- other historical weather variables
- climate model data
- demographic controls
- natural controls
- quality of life measure
- MSA assignment

We detail the data assembly process in each of the categories listed above in the following sections.

2 Historical temperature

Our source of historical temperature data is the Schlenker-Roberts daily weather data for the contiguous United States,¹ available from 1950–2005. They construct these data by combining high resolution *monthly* average PRISM temperature data (discussed more in detail in section 3) with *daily* temperature data obtained from weather stations, which are at a much lower spatial resolution. They regress the monthly average PRISM temperature data on monthly average weather station data, and then interpolate the daily temperature records at each weather station to obtain daily records at each PRISM grid point. For more details regarding the construction of these data, see their article “Nonlinear Effects of Weather on Corn Yields”.²

The data are set on the same fine grid as the PRISM data: 2.5 arcmin \times 2.5 arcmin (roughly 4 km \times 4 km) grid.

2.1 Binning temperature

Starting with the daily temperature data, we calculate, for each gridpoint-year, the number of days with an average temperature (mean of minimum and maximum temperature) falling in each of 222 temperature bins. The number of days in each bin is calculated as follows:

$$\sum_{day=1}^{365} 1\{avgTemp_{day} \in range(bin)\}$$

We define 220 bins ranging from -50 °C to 60 °C, with two more bins for days below -50 °C or above 60 °C, for a total of 222 bins; see table A2.1 for more detail.

¹ Schlenker and Roberts (2009). “Nonlinear Temperature Effects indicate Severe Damages to U.S. Crop Yields under Climate Change,” Proceedings of the National Academy of Sciences, 106(37), September 15 2009, p.15594-15598.

² Schlenker and Roberts. “Nonlinear Effects of Weather on Corn Yields,” *Review of Agricultural Economics*, 28(3), Fall 2006, p. 391–398.

Bin	Temperature Range, °C
bb001	$(-\infty, -50.0)$
bb002	$[-50.0, -49.5)$
bb003	$[-49.5, -49.0)$
...	...
bb100	$[-1.0, -0.5)$
bb101	$[-0.5, 0.0)$
bb102	$[0.0, 0.5)$
bb103	$[0.5, 1.0)$
...	...
bb220	$[59.0, 59.5)$
bb221	$[59.5, 60.0)$
bb222	$[60.0, \infty)$

Table A2.1: Temperature ranges covered by each bin.

Because there are no days with temperatures lower than -50 °C or higher than 60 °C, we can simply convert daily average temperature to its corresponding bin using:

$$bin = \text{floor}(2 \times \text{avgTemp}) + 102$$

The raw bins add up to slightly over 365 per year due to leap years, so we rescale the bins to add up to exactly 365 over a year. We perform this gridpoint-year binning for each year between 1970 and 1999, inclusive. As temperature and other weather indicators are generally averaged over 30-year “normals” (more specifically, the PRISM data we use follows the 1970–1999 normals), we generate normals of our binned data by averaging over these same 30 years.

2.2 Mapping binned temperature to PUMAs

We map the binned daily temperature data to PUMAs using one of the following two methods, depending on the size of the PUMA.

We use PUMA boundaries available in shapefile format from IPUMS USA.³ For each PUMA that has at least one PRISM grid point within its boundary, we match the PUMA to the binned temperature data for their gridpoints and average across the PUMA, weighting all grid points equally.

For those few (33) very small PUMAs with that contain no grid points, we match the PUMA to the binned temperature data for the four (or fewer) gridpoints surrounding the PUMA centroid, and take the inverse-distance-squared weighted average, using great-circle distances. PUMA centroids are calculated using ArcMap from PUMA boundaries.

Most PUMAs are matched to four grid points. However, for some coastal PUMAs, the corresponding over-water grid points do not have data, so they are matched to three or fewer gridpoints. The following two tables summarize the mapping from PRISM grid points to PUMAs:

³ <https://usa.ipums.org/usa/volii/2000pumas.shtml>

Count of PUMAs	
with grids inside	2,024
without grids inside	33

Count of grid points	
avg. per PUMA	228.9
falling inside a PUMA	470,773
not falling inside a PUMA	386
total	471,159

3 Other historical weather data

We assemble data on dew point, relative humidity, precipitation, and sunshine at the PUMA level.

Dew point, relative humidity, and precipitation come directly from, or are derived from, PRISM data. PRISM, or the Parameter-elevation Regressions on Independent Slopes Model climate mapping system is a “unique knowledge-based system that uses point measurements of precipitation, temperature, and other climatic factors to produce continuous, digital grid estimates of monthly, yearly, and event-based climatic parameters.”⁴ We used PRISM monthly average temperature, precipitation, and dew point data from the years 1970–1999. The data are available on a high-resolution grid of 2.5 arcmin × 2.5 arcmin (roughly 4 km × 4 km).

3.1 Dew point

The PRISM data give average dew point in each month for each PRISM grid point. We average over all 30 years to obtain the average dew point, for each grid point, in each calendar month, 1970–1999, and map these gridded data to PUMAs as described in section 2.2.

3.2 Relative humidity

We calculate relative humidity from PRISM temperature and dew point data, following conversions given by the equations below. We first convert temperature t and dew point d to degrees Celsius. Actual vapor pressure e and saturated vapor pressure e_s are calculated as follows:

$$e = 6.112 \cdot \exp\left(\frac{17.67d}{d + 243.5}\right)$$

$$e_s = 6.112 \cdot \exp\left(\frac{17.67t}{t + 243.5}\right)$$

Then relative humidity is given by:

$$rh = 100 \frac{e}{e_s}$$

One problem is that the equations defining relative humidity as a function of temperature and dew point are nonlinear, so true average relative humidity is likely not exactly the same as relative humidity calculated from average temperature and average dew point. Because dew

⁴ <http://www.prism.oregonstate.edu/https://usa.ipums.org/usa/volii/2000pumas.shtml>

point and temperature are not extremely variable within a given month at a given location, this is unlikely to significantly change the relative humidity calculation.

These gridded relative humidity data are then mapped to PUMAs as described in section 2.2.

3.3 Precipitation

We import gridded precipitation data from PRISM, which give average precipitation in each month over the periods 1970–1999, for each PRISM grid point. We average to obtain the mean precipitation for each grid point, in each calendar month, 1970–1999. These gridded data are then mapped to PUMAs as described in section 2.2.

3.4 Sunshine

Because high resolution gridded sunshine or cloud cover data are not available from PRISM or other sources, we use weather station data from the National Climactic Data Center.⁵ This data set contains data on the amount of time the sun is not obscured by clouds, as a percent of possible daylight hours, by month-of-year, for 174 weather stations across the United States; this includes some stations in Alaska, Hawaii, Puerto Rico, and other US territories that we do not use. We use only the 156 stations in the contiguous United States, mapped in figure A2.1. The data set also lists the total number of years of sunshine observations for each station, which ranges from only a few years to over 100. These data do not give the longitude or latitude of each weather station, but rather give us a unique identifier for each station. We look up each station in the *North American Master Station Database*⁶ to obtain the longitude and latitude.



Figure A2.1: Weather stations used for construction of sunshine dataset.

We then use inverse distance squared weighting to interpolate the weather station data to the PUMA level. The closest four weather stations, appropriately weighted, are used for the interpolation to each PUMA centroid. We use the great circle distances because the distances

⁵ <http://www.ncdc.noaa.gov/oa/climate/online/ccd/pctpos.txt>

⁶ <http://www.weathergraphics.com/identifiers/master-station.dat>

between PUMA centroid and nearest four weather stations can be potentially large. Interpolated values of sunshine are not sensitive to the number of surrounding weather stations used in the interpolation. These data are mapped in figure A2.2.

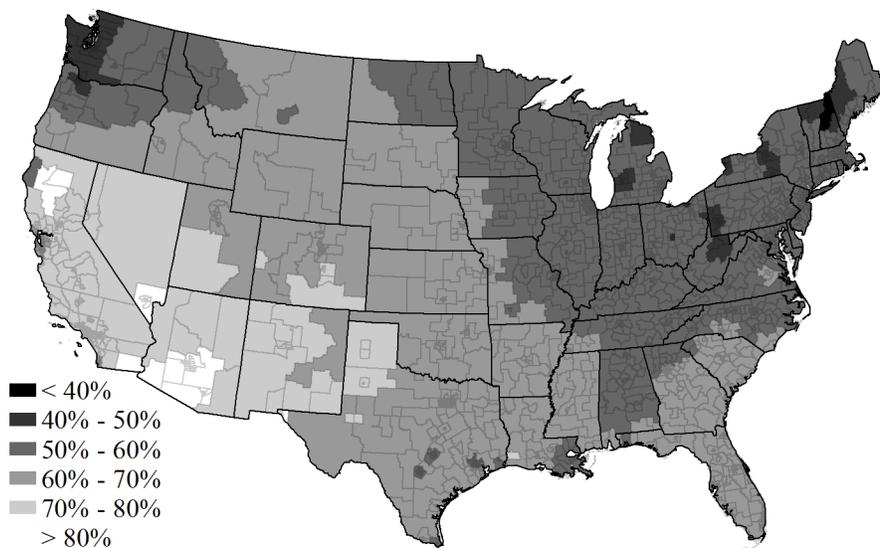


Figure A2.2: Annual sunshine, in percent of possible daylight hours.

4 Climate model data

All climate change scenario data are from the third release of the Community Climate System Model (CCSM 3.0), available from the Earth System Grid Gateway at the National Center for Atmospheric Research.⁷ These data were also used for the 2007 Fourth Assessment Report of the United Nations Intergovernmental Panel on Climate Change.

We use average daily and monthly data from five runs (a–e) of the *Climate of the 20th Century experiment* from the years 1970–1999, five runs (a–e) of the *A2 scenario* from 2000–2099, and the only available run (a) of the *A1FI scenario* from 2000–2099. The data are available on a 1.4 degree \times 1.4 degree resolution grid.

4.1 Daily data

The climate model provides daily data for all variables; we use daily average surface temperature data (CCSM field name *TS*) to construct binned temperature, similarly to the construction of binned data from the Schlenker-Roberts daily weather data described in section 2. For each decade, we calculate the average number of days that fall in each bin as described in section 2.1.

4.2 Monthly data

We use monthly average surface temperature, relative humidity, surface pressure, convective precipitation rate, and large-scale (stable) precipitation rate (field names *TS*, *RELHUM*, *PS*, *PRECC*, and *PRECL*, respectively), calculating decadal averages of each variable in each calendar month. Convective and large-scale precipitation rates are added together to

⁷ <http://earthsystemgrid.org>

obtain total precipitation.

4.3 Mapping to PUMAs

These monthly average and binned daily average data are then mapped to PUMAs. Because the CCSM grids are significantly bigger than the PRISM grids described in 2, there are many PUMAs that do not have even one CCSM grid point falling within its boundary. Thus, all PUMAs are assigned the inverse-distance-squared weighted value of the four grid points surrounding the centroid of the PUMA, as for the 33 PUMAs interpolated using the alternate method in section 2.2. We use the true great-circle distances, as the distances between PUMA centroid and surrounding grid point can be large, up to 100 km.

We calculate PUMA-level changes in the temperature distribution as the bin-by-bin difference between 1970–1999 *20CE* climate model data and the *A2* or *A1FI* climate model data. However, this yields a strange result for some PUMAs: if the climate model does not predict the actual temperature distribution correctly, it is possible that the projected decrease in the number of days in some bin will be larger than the number of days that we actually observe in that bin originally. This yields a predicted future temperature distribution with a few bins that have a negative number of days.

Fortunately, these errors are small. We calculate the number of bins with a negative value (of less than -0.1 days) for 2070–2099 predicted temperatures under the *A2* scenario; across all PUMAs, this has a mean of 11.9 bins. As a percentage of binned days, the number of negative predicted binned days is likewise small, with a mean of 3.9%.

5 Demographic controls

We collect various demographic variables, including population, population change, schooling, age, percent married, percent veteran, percent in each minority group, and other variables. We calculate PUMA-level statistics for various variables over all individuals in each PUMA (not just workers). We use data from IPUMS.

Also, we calculate a population-weighted version of population density at the PUMA level, and impute population change from 1990–2000, as described in the following sections.

5.1 Population density data

We construct a population density measure that more accurately reflects the actual population density experienced by the population, relative to a simple *population/area* (that is, *area-weighted*) measure over an entire county PUMA. We acquire population and (land) area data by census tract from the 2000 Census.⁸ Census tracts are subdivided much more finely than PUMAs; there are generally only between 2000–8000 people in each tract.

⁸ <http://dataferrett.census.gov/>

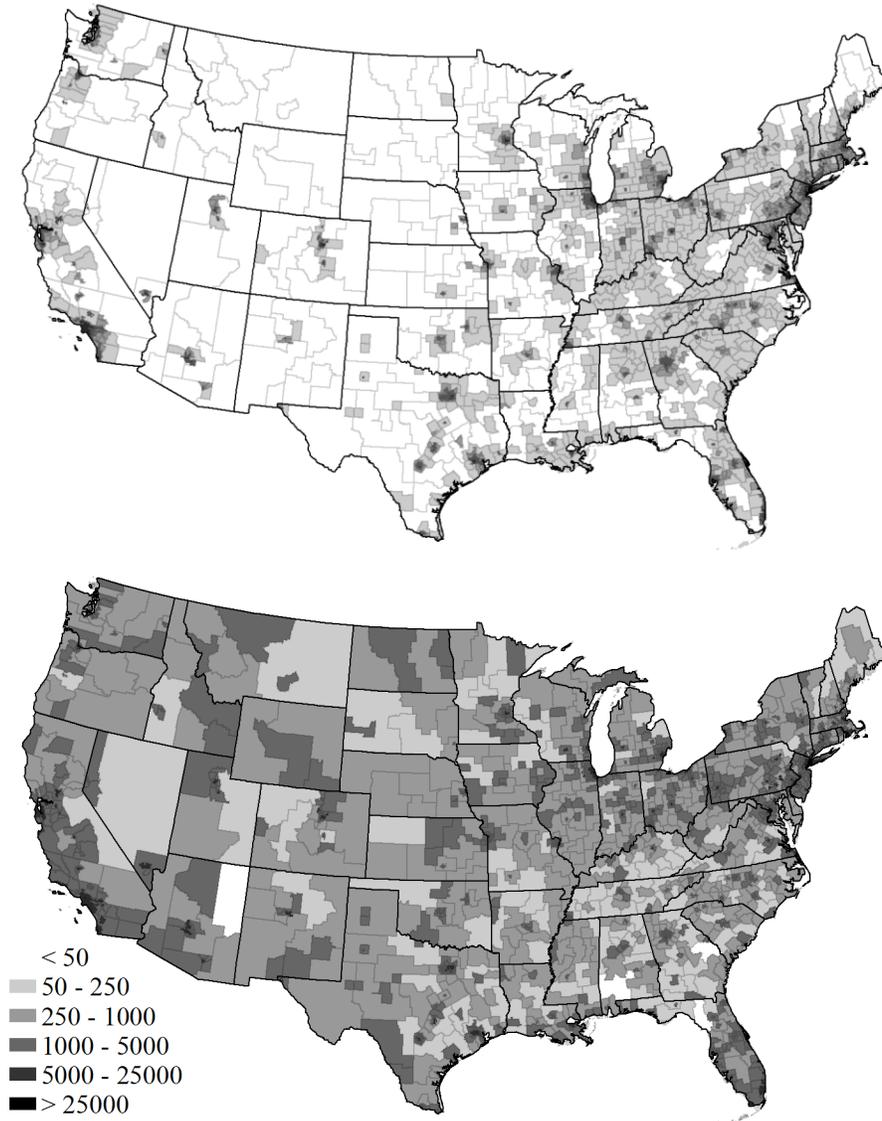


Figure A2.3: Comparison of unweighted and weighted population density by PUMA, in persons per square mile.

Our new population density data on PUMA level are computed by first calculating simple population density ($population/area$) at the census tract level. Because census tracts do not cross county lines, computing this average would be simple for counties. However, census tracts do sometimes cross PUMA boundaries; thus, we use MABLE/Geocorr2K correspondence files to assign the correct areas and populations of the census tracts to each PUMA. Thus, for each PUMA p , the population weighted density is computed as follows:

$$weightedDensity_p = \frac{\sum_{t \in p} density_t \cdot population_t \cdot fract(t \in p)}{\sum_{t \in p} population_t \cdot fract(t \in p)},$$

where, for each tract t and each PUMA p , MABLE/Geocorr2K gives the allocation factor:

$$\text{fract}(t \in p) = \frac{\text{total population in } t \text{ and in } p}{\text{total population in } t}$$

The weighted and unweighted PUMA densities are mapped in figure A2.3.

5.2 Population change data

We are interested in the PUMA-level population change between the 1990 census and 2000 census. However, because PUMA definitions change every census, we cannot directly use the 1990 PUMA population data, as the PUMAs do not line up. So, we impute 1990 population in our 2000 PUMAs using 1990 census tract level data, as follows.

In ArcMap, we intersect the 1990 tract shapefile with the 2000 PUMA shapefile; this splits each tract into one or more pieces, each one lying inside a different PUMA. Then, for each tract slice, we calculate the area, and compute the fraction of the whole tract area lying inside each PUMA. Then, we attribute that same fraction of the 1990 tract population to that PUMA, and add up over tracts to get the whole PUMA population.

The true 1990 population, not including AK and HI, is 247,051,601, while the total population we calculate using this method is 246,956,773. The discrepancy comes from a few hundred tracts that are not included in the PUMA shapefiles, and thus were not attributable to any PUMAs. These are tracts ending .99, which according to the US Census Bureau⁹ is reserved for tracts consisting only of “crews-of-vessels populations.” We are unfortunately unable to assign this small sub-population properly to PUMAs, resulting in a missing population on the order of 0.038%.

For a consistency check, we repeat the above procedure to impute 2000 PUMA populations from 2000 census tracts. For the 2000 census, the crews-of-vessels populations are part of the related census tracts, so we do not have the same problem as above; the true 48-state population is 279,583,437, which is exactly the total population we calculated. However, the population we assigned to each PUMA is not exactly correct, and thus will not match the true PUMA populations, as it is imputed using this imperfect area-weighting method.

We approximate population change by taking the difference between imputed PUMA population in 2000 and imputed PUMA population in 1990.

6 Natural controls

We construct data on distance to coasts and mountainousness at the PUMA level from raw elevation data and lake and ocean shapefiles, as detailed in the following sections.

Specifically, we construct

- average elevation (in meters)
- average slope (in percent)
- distance to ocean (in miles)
- indicator for whether PUMA is on the ocean
- distance to a great lake (in miles)
- indicator for whether PUMA is on a great lake

6.1 Elevation data

We use high-resolution elevation data from the *Global 30 Arc Second Elevation Data*

⁹ http://www.census.gov/geo/www/cob/tr_metadata.html

(GTOPO30) digital elevation model (DEM) available from the the United States Geological Survey.¹⁰ These data are set on a high resolution grid of 30 arcsec \times 30 arcsec, or roughly 1 km \times 1 km. We map the gridded elevation data to PUMAs by averaging the value of all grid points falling within the boundary of each PUMA. Due to the high resolution of the data, every PUMA has at least one grid falling within its boundary. Elevation data are mapped in figure A2.4.

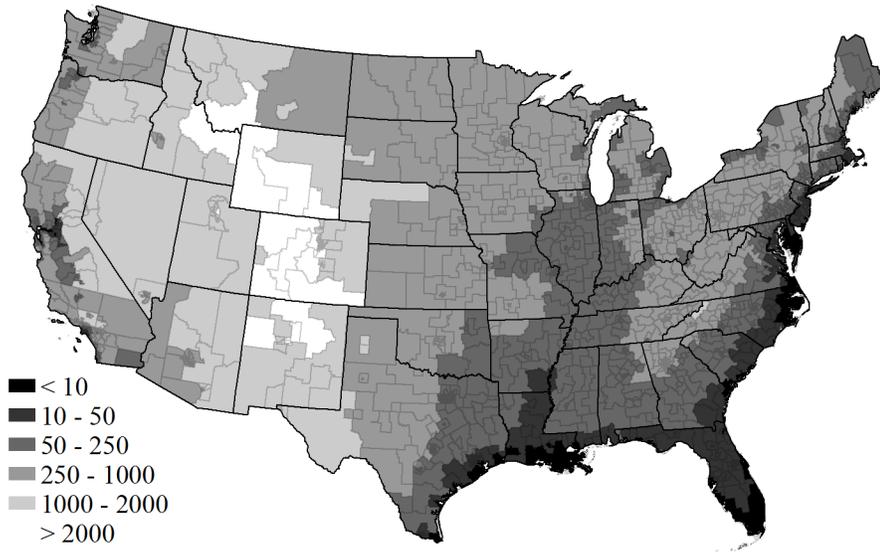


Figure A2.4: Average elevation by PUMA, in meters.

6.2 Land slope data

We construct a land slope dataset using the DEMs described in section 6.1. We calculate slope using the average maximum technique, where the slope at each grid point is the maximum rate of change of elevation from that grid point to its eight neighbors.¹¹ ¹² Consider a 3 \times 3 neighborhood of cells, where each cell is identified as follows:

a	b	c
d	e	f
g	h	i

Then the slope at gridpoint e, in radians, is given by:

$$slope = \arctan \sqrt{\left(\frac{\partial z}{\partial x}\right)^2 + \left(\frac{\partial z}{\partial y}\right)^2},$$

where

¹⁰ http://eros.usgs.gov/#/Find_Data/Products_and_Data_Available/gtopo30_info

¹¹ ArcToolbox help, section "How Slope Works," ESRI.

<http://webhelp.esri.com/arcgisDesktop/9.3/index.cfm?TopicName=How%20Slope%20works>

¹² Burrough, P. A. and McDonell, R.A., 1998. *Principles of Geographical Information Systems* (Oxford University Press, New York), p. 190.

$$\frac{\partial z}{\partial x} = \frac{(c + 2f + i) - (a + 2d + g)}{8 \cdot cellSize_x}$$

$$\frac{\partial z}{\partial y} = \frac{(g + 2h + i) - (a + 2b + c)}{8 \cdot cellSize_y}$$

As in section 6.1 above, we map the gridded slope data to PUMAs by averaging the slope of all grid points falling within the boundary of each PUMA. Slope data are mapped in figure A2.5.

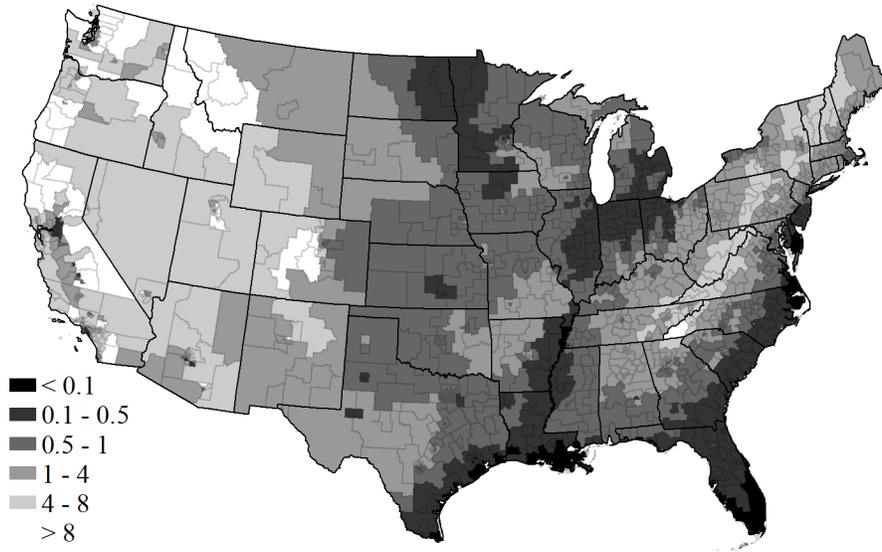


Figure A2.5: Average slope by PUMA, in percent

6.3 Coast data

For each geographic level, we construct the following variables:

1. Distance from centroid to oceanic coast
2. Indicator variable: $1_{\{onOceanCoast\}}$
3. Distance from centroid to Great Lakes coast
4. Indicator variable: $1_{\{onGreatLakesCoast\}}$

To obtain the distance variables (1) and (3), we first export the set of points defining the oceanic and Great Lakes from shapefiles available from Esri, creator of ArcMap and the ArcGIS suite. We then calculate the minimum distance from each PUMA centroid to any oceanic or Great Lakes coast point.

Indicator variables for a PUMA being on an oceanic or Great Lakes coast were calculated similarly, with the difference that each PUMA is first assigned the *minimum* distance between any point defining its border and any point on a coast. If this minimum distance is zero (or nearly so, allowing for some small errors in the shapefiles), the coastal indicators are assigned a value of one, and otherwise, a value of zero.

7 Quality of life data

Quality of life data at PUMA level come from Albouy and Lue (forthcoming). We use both commuting-adjusted and non-adjusted QOL, data on average QOL and QOL by education

level (high school or college), and price and wage differentials. Construction of these data are detailed in Albouy and Lue (forthcoming).

8 MSA assignment

We use MABLE/Geocorr2K¹³ to map each PUMA to its corresponding consolidated metropolitan statistical area (CMSA). Each PUMA is matched to a single CMSA, but CMSAs contain more than one PUMA; the fraction of the CMSA population contained within each PUMA is also given by MABLE/Geocorr2K.

¹³ <http://mcdc2.missouri.edu/websas/geocorr2k.html>

Appendix 3: Concavity of the utility function

The specifications allowing for unobserved preference heterogeneity use households' first order condition to identify their MWTP for climate amenities at their location, following the approach of Bajari and Benkard (2005) under an assumption that the choice set is continuous. This approach cannot identify, however, the shape of households' utility functions as one moves away from the characteristics associated with their location. Because climate change is non-marginal, this shape affects the estimates of amenity loss associated with climate change. To be conservative, all results presented in the main text of the paper assume that households' utility functions are linear rather than concave in local characteristics.

The linearity assumption may be problematic, however, because it can lead to violations of households' utility maximization should the second order condition not hold. If the hedonic price function is not convex, then households satisfying the first order condition with linear utility functions will be locally minimizing rather than maximizing their utility at any location where the price function is locally concave. Further, even if a particular household is locally maximizing utility at a location where the price function is locally convex, concavities elsewhere in the price function may imply that the household would realize a higher utility at a different location.

In our application, we do find some concavities in the estimated price function, suggesting that violations of utility maximization may be a problem. To assess the empirical importance of this issue, we calculate the amount of utility function concavity necessary, at each location, to satisfy utility maximization. To do so, we use a utility function parameterized per equation (A3.1) below to denote the utility that a household located at i would obtain from residing at location j :

$$u_j^i = \beta^i Z_j - \alpha^i \sqrt{\sum_{k=1}^{K'} \left(\frac{Z_{ik} - Z_{jk}}{\sigma_k} \right)^2} - q_j + \xi_j \quad (\text{A3.1})$$

The quadratic term in (A3.1) allows the MWTP for each characteristic k to decrease as the difference between the amount of k provided at locations i and j increases.¹ All characteristics are scaled by σ_k , their standard deviation over all locations, so that their magnitudes are comparable.²

The household's first order condition yields estimates of the β per the local linear regressions discussed in the main text. These regressions also yield estimates of the ξ_j . Given these estimates, for each location i we loop through all other locations j and find the smallest positive value of α^i such that $u_i^i \geq u_j^i \forall j$. The resulting estimates of α^i denote the minimum concavity necessary so that all households are maximizing their utility by residing at their present location.

In our baseline specification (restricted 7th degree cubic spline with full controls, and no fixed effects), the mean value of α^i obtained via this procedure is 0.0012. When calculating losses from climate change, allowing for this degree of concavity yields an additional welfare loss of 0.36% of income.

¹ The quadratic terms only involve characteristics $1, \dots, K$ for which preferences are permitted to be heterogeneous. For all models presented in the main text, these characteristics are those corresponding to the basis functions for the linear or cubic splines over the temperature bins.

² In a more general specification, the α^i coefficients on the quadratic term could vary across characteristics. For example, in a model with both HDD and CDD, utility maximization could be achieved with: (i) high concavity on HDD but none on CDD; (ii) high concavity on CDD but none on HDD; or (iii) moderate concavity on both HDD and CDD. For simplicity, we assume that the each household's concavity coefficient α^i is the same for all characteristics in the model.

Appendix 4: Set identification of preferences

The estimated heterogeneous preferences presented in the text (and the climate change welfare impacts derived from them) assume point identification, following Bajari and Kahn (2005). That is, at each location i , the estimates assume that households' MWTP for a given characteristic k is given by the gradient of the price function at i , β_k^i , per households' first order condition. Point identification, however, requires that the set of available locations is complete in characteristics space (itself a subset of \mathbb{R}^K), a condition that cannot hold in a finite choice set such as our sample of PUMAs. Thus, in reality, preferences are only set identified. However, if the number of choices (in our data, the 2,057 PUMAs) is sufficiently large, the identified sets may be small enough that estimates assuming point identification are close approximations to the identified set. In particular, the distribution of point identified preferences across PUMAs may closely match the distribution of set identified preferences. This appendix examines the extent to which this approximation holds in this setting.

At each location i , finding the identified set of preferences requires finding the set of β^i 's such that $u_i^i \geq u_j^i \forall j$. We use a utility function that is concave in characteristics and given by equation (A4.1) below (copied from equation (A3.1) in appendix 3):

$$u_j^i = \beta^i Z_j - \alpha^i \sqrt{\sum_{k=1}^K \left(\frac{Z_{ik} - Z_{jk}}{\sigma_k} \right)^2} - q_j + \xi_j \quad (\text{A4.1})$$

An allowance for concavity is required to ensure that the identified set is not null. Given values for the characteristics Z and quality of life "prices" q , and given estimates of the unobservables ξ_j and concavity parameters α^i from the local linear regressions and the procedure discussed in appendix 3, respectively, equation (A4.1) can be used to find the identified set of β^i at each location.¹ To do so, we use apply the Bayesian estimation procedure discussed in Bajari

¹ As discussed in appendix 3, the α^i are set to ensure that utility maximization holds exactly at each location, given point identified estimates of the β^i . As evidenced by the fact that the set identified preferences are not singletons, these α^i do not imply that other values of β^i are inconsistent with utility maximization. Increases in α^i do imply increases in

and Benkard (2005) and described in detail in Benkard and Bajari (2004).² This procedure assumes that the distribution of preferences, within and across locations, is point identified and uses Gibbs sampling to obtain this distribution (an analytical characterization of these sets is very difficult because they are high-dimensional).

The Gibbs sampler requires a prior on the preference distribution. Following Benkard and Bajari (2004), we assume that preferences are distributed uniformly over a rectangular region \mathcal{B} . We must place bounds on this region so that, at locations with extreme characteristics, preferences are not unbounded (for example, at the hottest PUMA in the data, residents could in theory have an unbounded taste for heat). For each characteristic k over which preferences may be heterogeneous, we bound the set \mathcal{B} in the k 'th dimension by taking the support of β_k from the point identified estimates and doubling its width. Thus, for example, if the point identified MWTPs for characteristic k range from -0.1% to -0.2% of income, our prior for the Gibbs sampler is that preferences for k are uniformly distributed from -0.05% to -0.25% of income.

Given our prior, the Gibbs algorithm proceeds one location i at a time. A summary of the algorithm's steps in each iteration is given below. For details, see Benkard and Bajari (2004).

1. At a given iteration s , start with characteristic $k = 1$. Find the upper and lower bounds of β_k^i such that $u_i^i \geq u_j^i \forall j$, given the β_{-k}^i from iteration $s-1$ (use the point-identified β_{-k}^i if $s = 0$). Denote these bounds by $\bar{\beta}_k^i$ and $\underline{\beta}_k^i$. If location i has the minimum or maximum value of X_k in the choice set, then the lower or upper bound, respectively, will be given by the relevant bound of \mathcal{B} .
2. The posterior distribution of β_k^i for iteration s is uniform on $[\underline{\beta}_k^i, \bar{\beta}_k^i]$. Randomly draw a β_k^i from this distribution.

the size of the identified set of β^i ; however, this sensitivity is not strong. Doubling the values of α^i from those calculated in appendix 3 yields only a minor change in the kernel density plots discussed below.

² Benkard, C. Lanier and Patrick Bajari (2004) "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach," NBER working paper 10278.

3. Repeat steps 1 and 2 for each characteristic $k \in 2, \dots, K'$ (K' denotes the number of characteristics over which preferences are permitted to be heterogeneous). When finding the bounds $\bar{\beta}_k^i$ and $\underline{\beta}_k^i$, use β_l^i from iteration s rather than $s-1$ if $l < k$.

We use 1,500 draws in our algorithm and discard the first 500 as a “burn in” period. The remaining 1,000 draws at each location i then give a numerical estimate of the distribution of preferences at i . Aggregating the draws over all locations yields a numerical estimate of the distribution of preferences over the population. Figures A4.1 and A4.2 below show kernel density plots of the MWTPs for a 40 degree day and an 80 degree day (both relative to 65°F) for the 7th degree cubic spline specification that includes demographic controls. The primary difference between the point and set identified MWTP estimates is that set identification yields long, thin tails of extreme preferences that are driven by locations at or near the extremes of the characteristics space.³ Otherwise, these two figures show strong agreement between the preference distributions implied by point identification and set identification, validating the use of the point identified estimates in the main text.

As noted in the text, we only allow for preference heterogeneity over the characteristics associated with temperatures and enforce preference homogeneity for the controls. We do so because the close agreement between the set and point identified preference distributions breaks down when we allow for heterogeneous preferences for the control characteristics. Figures A4.3 and A4.4 illustrate kernel density MWTP plots when preferences for the geographic and “other weather” controls are permitted to be heterogeneous. When preferences for this many characteristics are permitted to vary, the identified sets become too large to allow precise inference.

³ Note that, even though the space \mathcal{B} has bounds given by a doubling of the support of the set of point identified preferences, the set identified MWTPs for a given temperature can have bounds given by more than a doubling of the support of point identified MWTPs. This outcome is due to the fact that the MWTP for a given temperature is given by a linear combination of multiple characteristics (seven of them in the cubic spline models).

Figure A4.1: Kernel density plots of the MWTP for a 40°F day

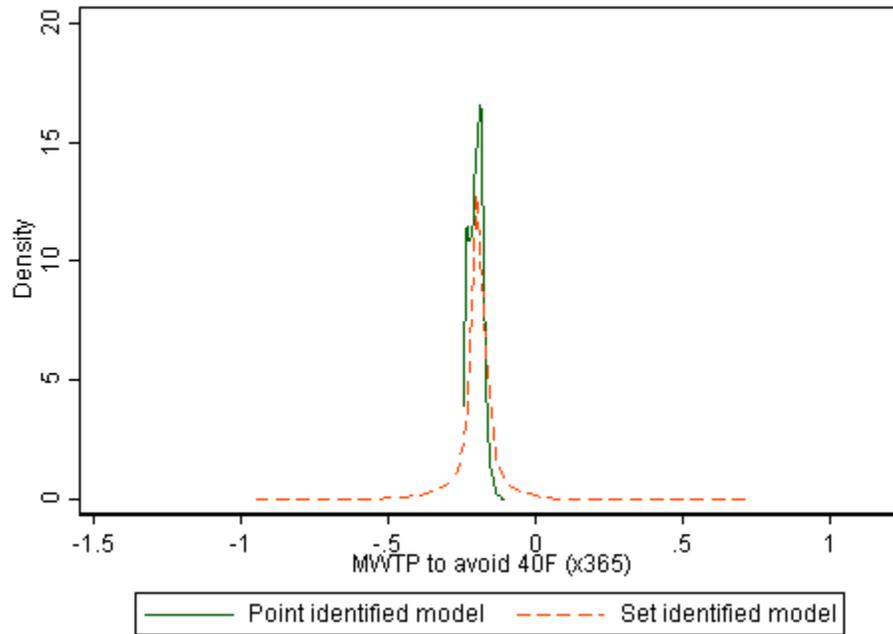
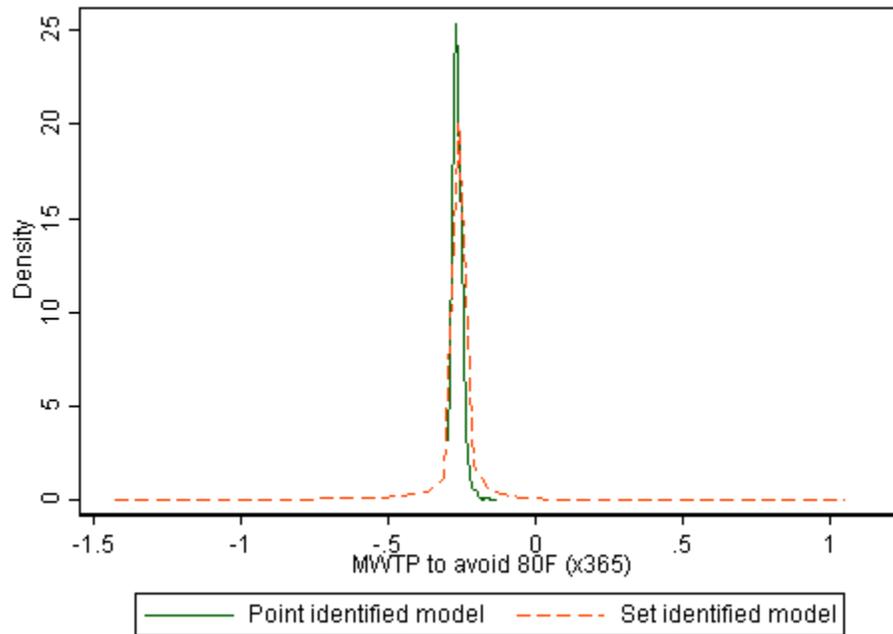


Figure A4.2: Kernel density plots of the MWTP for an 80°F day



Notes: Both figures use the 7th-degree restricted cubic spline model with demographic controls. The only random coefficients are those on the temperature basis functions.

Figure A4.3: Kernel density plots of the MWTP for a 40°F day
Estimates allow preference heterogeneity for geographic and “other weather” controls

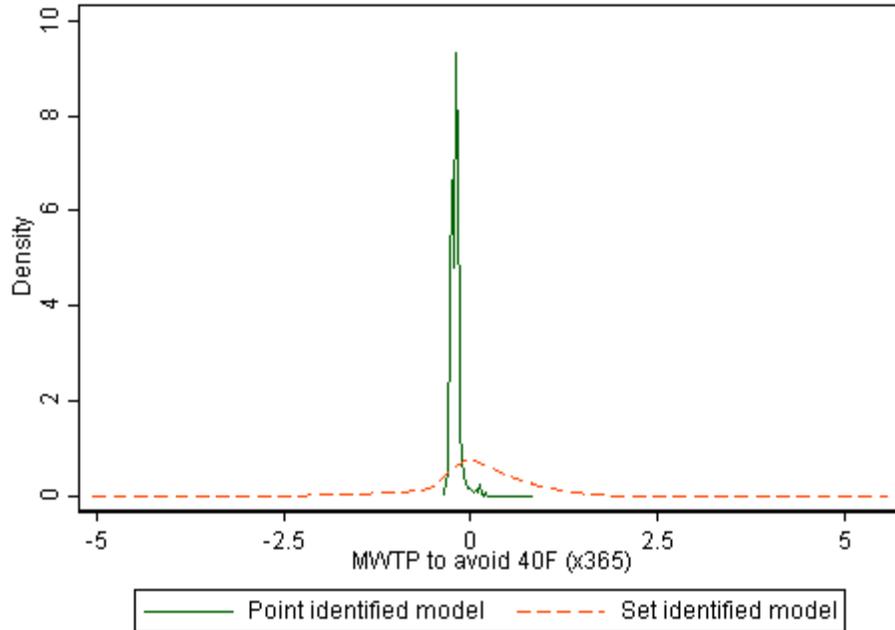
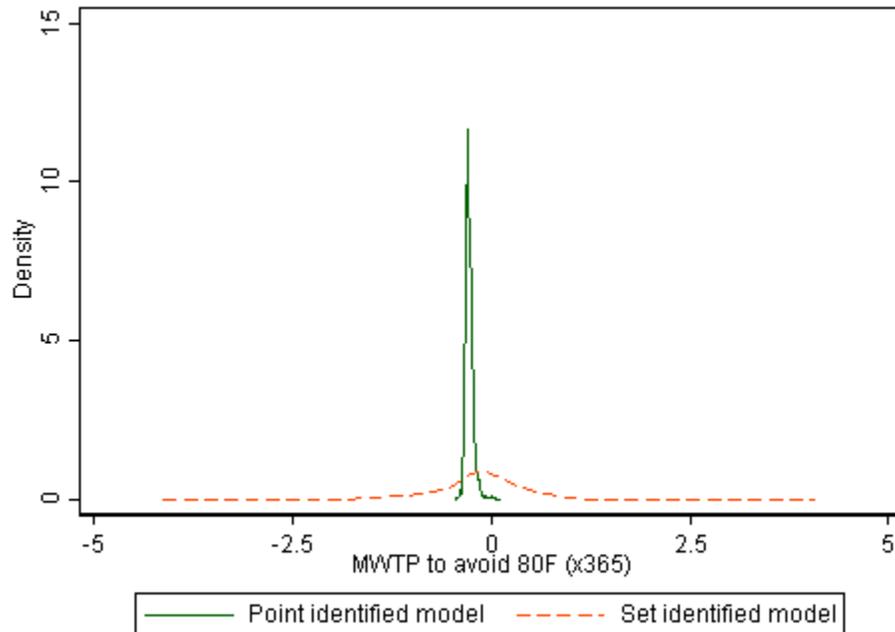


Figure A4.4: Kernel density plots of the MWTP for an 80°F day
Estimates allow preference heterogeneity for geographic and “other weather” controls



Notes: Both figures use the 7th-degree restricted cubic spline model with demographic controls. Random coefficients are used on the temperature basis functions, geographic controls, and other weather controls.