

NBER WORKING PAPER SERIES

DOES THE ENVIRONMENT STILL MATTER? DAILY TEMPERATURE AND
INCOME IN THE UNITED STATES

Tatyana Deryugina
Solomon M. Hsiang

Working Paper 20750
<http://www.nber.org/papers/w20750>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2014

We thank Max Auffhammer, Marshall Burke, Tamma Carleton, Melissa Dell, Michael Greenstone, Amir Jina, Edward Miguel, Matthew Neidell, Julian Reif, Wolfram Schlenker, Richard Schmalensee, James Stock, Anna Tompsett, Reed Walker, and seminar participants at MIT and UIUC for discussions and comments. We thank DJ Rasmussen and Michael Delgado for research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Tatyana Deryugina and Solomon M. Hsiang. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Does the Environment Still Matter? Daily Temperature and Income in the United States
Tatyana Deryugina and Solomon M. Hsiang
NBER Working Paper No. 20750
December 2014
JEL No. N52,O1,O4,Q51,Q54,R11

ABSTRACT

It is widely hypothesized that incomes in wealthy countries are insulated from environmental conditions because individuals have the resources needed to adapt to their environment. We test this idea in the wealthiest economy in human history. Using within-county variation in weather, we estimate the effect of daily temperature on annual income in United States counties over a 40-year period. We find that this single environmental parameter continues to play a large role in overall economic performance: productivity of individual days declines roughly 1.7% for each 1°C (1.8°F) increase in daily average temperature above 15°C (59°F). A weekday above 30°C (86°F) costs an average county \$20 per person. Hot weekends have little effect. These estimates are net of many forms of adaptation, such as factor reallocation, defensive investments, transfers, and price changes. Because the effect of temperature has not changed since 1969, we infer that recent uptake or innovation in adaptation measures have been limited. The non-linearity of the effect on different components of income suggest that temperature matters because it reduces the productivity of the economy's basic elements, such as workers and crops. If counties could choose daily temperatures to maximize output, rather than accepting their geographically-determined endowment, we estimate that annual income growth would rise by 1.7 percentage points. Applying our estimates to a distribution of "business as usual" climate change projections indicates that warmer daily temperatures will lower annual growth by 0.06-0.16 percentage points in the United States unless populations engage in new forms of adaptation.

Tatyana Deryugina
Department of Finance
University of Illinois, Urbana-Champaign
deryugin@illinois.edu

Solomon M. Hsiang
Goldman School of Public Policy
University of California, Berkeley
2607 Hearst Avenue
Berkeley, CA 94720-7320
and NBER
shsiang@berkeley.edu

1 Introduction

Since the invention of agriculture, one narrative of economic progress has been a gradual decoupling of humanity’s economic well-being from an unpredictable environment. As economies have progressed from hunting and gathering groups to the sophisticated systems we participate in today, the welfare of individuals within these economies seems to depend less and less on the specific environmental conditions in which the economy is embedded. This profound transition has been driven by technological innovations that allow individuals to transform naturally occurring capital—such as minerals, plants, or sunlight—into a human-engineered capital stock that is more robust to environmental conditions (Hartwick (1977), Solow (2012)).

Here we examine how far this process has taken us. We test whether local environmental conditions continue to influence economic well-being in the wealthiest economy to ever exist. We exploit natural random variation in a single environmental parameter, daily temperature, to identify its effect on per capita income in modern US counties. Controlling for annual fluctuations, county fixed effects, and the distribution of daily rainfall within each county-year, we use a difference-in-differences approach to recover the non-linear effect of daily temperature on income. We obtain estimates that are both economically significant and statistically precise: the average productivity of individual days declines roughly linearly by 1.7% for each 1°C (1.8°F) increase in daily average temperature beyond 15°C (59°F). This response is driven by warm weekdays and is not attenuated in the most recent and wealthiest years of our sample—it is easily obtained when we restrict our analysis to the 21st century.

Determining if wealthy economies are vulnerable to the impact of high temperatures is a particularly critical question in the economics of global climate change. Should wealthy economies be affected by temperature, then a much larger fraction of the global economy may be disturbed by climate change than previously thought. For example, in an early assessment Schelling (1992) stated,

“I conclude that in the United States, and probably Japan, Western Europe, and other developed countries, the impact [of climate change] on economic output will be negligible and unlikely to be noticed.”

Discussions and analysis by Poterba (1993), Stern (2006), Nordhaus (2008), and Tol (2009), among others, make similar assumptions or arrive at similar conclusions. This notion has been reinforced by reduced-form, linear regressions in Dell, Jones, and Olken (2012), who document that annual economic growth rates in poor countries are negatively correlated with annual variations in each country’s temperature but no statistically significant correlation appears to exist for richer countries.

It is generally thought that output in wealthy economies is less affected by climatic conditions because these populations are better equipped to adapt to their environment (Kahn (2005)). For example, farmers may alter what crops they plant to maximize profits conditional on their climate (Mendelsohn, Nordhaus, and Shaw (1994)), they may innovate to develop new varieties robust to

adverse climates (Olmstead and Rhode (2011)), or they may simply migrate away from deteriorating climates (Hornbeck (2012)). In other contexts, individuals may make defensive investments that minimize the impact of the climate, such as purchasing air conditioners (Barreca et al. (2013)). Finally, transfers or trade may be utilized to minimize the overall effect of climatic conditions on individual welfare (Deryugina (2014), Costinot, Donaldson, and Smith (2012)).

Yet a growing body of evidence suggests that even in wealthy countries some economic vulnerabilities remain, implying that adapting to all climatic conditions along all margins is too costly. For example, US crop yields continue to be highly sensitive to extreme heat, even after decades of investment in genetic development and infrastructure (Roberts and Schlenker (2011), Burke and Emerick (2013)), and measurable heat tolerance is only achieved at the cost of average yield reductions (Schlenker, Roberts, and Lobell (2013)). In non-agricultural contexts, the modern response of temperature-exposed labor supply (Graff Zivin and Neidell (2014)) is relatively unchanged from responses documented seventy years earlier (Mackworth (1947)); high temperatures continue to elicit costly personal conflicts even in wealthy populations (Hsiang, Burke, and Miguel (2013)); and adaptation to extreme climatic events like hurricanes remains only partial (Hsiang and Narita (2012)).

Here we examine whether, net of all potential adaptations, there remains an effect of temperature on income in the modern US (1969-2011). If relevant adaptive behaviors and investments are both pervasive and effective, then we expect local economic activity in US counties to be unaffected by temperature. However, if ambient temperature still matters, perhaps because the full range of adaptations needed to decouple the US economy from the climate is too costly or because certain impacts are unavoidable, then we expect variations in county-level income to reflect changes in local temperatures.

To identify the effect of temperature on income, we exploit annual within-county variations in the distribution of daily temperatures. We apply a difference-in-differences approach that allows us to estimate the marginal effect of a single day’s temperature on end-of-year income, conditional on temperatures experienced during the other 364 days of the year (Deschênes and Greenstone (2011)). This method does not rely on “shocks” to annual average temperatures but instead measures how income responds to relatively small distortions in the distribution of daily temperatures, i.e. unpredictable changes in local climate. Key to this approach is the fact that the sum of positive and negative disturbances to daily income manifests itself as a change in total income at the end of the year. If hot days generate systematically lower daily incomes, then counties that experience hotter-than-average weather during a given year will on average exhibit lower end-of-year incomes. Specifically, an additional hot day will lower annual income by that day’s income times the fractional loss due to the temperature on that day. We identify this marginal effect of temperature on daily income using the residual annual variations in both the count of days at each temperature and income per capita that remain after county fixed effects, year fixed effects, and auto-regressive terms have absorbed average cross-sectional, trending, and serial correlations in both variables.

We find that total personal income per capita is highest when 24-hour average temperatures

are between 9-15°C (48.2-59°F) and declines substantially on hotter days. This average effect is driven by lost income during warm and hot weekdays and is largely unchanged between 1969 and 2011. There is limited evidence of temporal or spatial displacement of income. Farm income loss would be larger absent crop insurance payouts, but other government transfers do not appear to meaningfully offset income losses. Because our income measures capture price changes, we conclude that changes in prices do not appear to ameliorate the effects of temperature.

Farm and non-farm incomes respond somewhat differently to temperatures, with the overall effect on total income appearing to be a weighted average of the two. The non-linear structure of these responses suggest that temperature matters because it reduces the productivity of basic components of economic activity: workers and crops. Overall, each additional warm day (24-27°C or 75.2-80.6°F) reduces an average county's total income per capita by roughly \$14.78, with non-farm income representing roughly \$3.03 of this loss. Each additional hot day (>30°C or >86°F) reduces total income per capita by \$20.56, equal to 28% of average daily income.¹ Integrating these effects over individuals suggests that daily temperature has an economically meaningful effect on economic production in the modern US.

Like many factors that affect production, there appears to be a well-defined optimal temperature for productive activities at the county level. However, the distribution of daily temperatures is a geographically-determined endowment that cannot generally be altered. To develop a sense for the overall economic importance of this fixed environmental factor, we engage in a thought experiment where we imagine that daily temperature *is* a choice variable that agents can freely optimize over to maximize profits. Specifically, agents in each county set each day's temperature to 12-15°C (53.6-59°F), the income-maximizing temperature. We then estimate how different income in the US would be in this scenario relative to historical experience, effectively computing the economic penalty borne by the US because counties cannot freely choose their daily temperature. Allowing this imaginary US economy to evolve from initial conditions defined by actual income in 1969, we find that the inability to control county temperatures lowers the annual growth rate of average income by 1.69 percentage points. Although this is only a partial equilibrium estimate, it suggests that we should carefully consider the economic importance of basic environmental conditions even in rich countries.

While we encourage caution when interpreting the above thought experiment, we think it may help inform an important debate in economics. It is easily observed in cross-section that higher temperatures around the world, both within and between countries, are associated with lower incomes on average. Some authors have argued that this association is causal, with geographic conditions playing a major role in patterns of economic development (e.g. Gallup, Sachs, and Mellinger (1999); Nordhaus (2006)). Other authors disagree, positing instead that temperature and other geographic measures have little effect on economies and the observed cross-sectional correlation is due to important omitted variables, such as social and political institutions, that are correlated with geographic conditions (e.g. Acemoglu, Johnson, and Robinson (2002); Easterly

¹0.08% of average annual income.

and Levine (2003)). We cannot settle this question in the global sense, but we believe that the above thought experiment provides the first estimate, albeit for a single country, of the economic burden imposed by geographically-determined temperature. Because our identification strategy only exploits variation in the distribution of daily temperatures within a county over time, it cannot be confounded by cross-sectional patterns of social or political institutions.

We conclude by simulating income growth in the US economy under a “business as usual” global climate change scenario (RCP 8.5). We simulate national economic performance using the 44 climate projections developed by Houser et al. (forthcoming), which, taken together, correspond to the probability distribution of the unknown global climate sensitivity (Weitzman (2009)). Under the median projection, we estimate that the US annual income growth rate will change by -0.12 percentage points relative to a pre-1990 climate (median estimate) due to higher daily temperatures, with the 10-90 centile range of effects spanning -0.15 to -0.09 percentage points per year.

The remainder of the paper is organized as follows. In Section 2, we develop a framework for the various ways in which temperature can affect economic production, allowing for the possibility that agents adapt through reallocation of resources, defensive investments, price adjustments, and transfers. In Section 3, we present our data. In Section 4, we describe our estimation strategy, including a derivation for how we recover daily responses to daily temperature when we do not have daily economic observations. In Section 5, we present our econometric estimates for the effect of daily temperature on income, along with several extensions. In Section 6, we implement simulations where counties can choose their temperatures and where climate change continues under “business as usual” assumptions. We discuss these findings and conclude in Section 7.

2 Framework

Consider an economy with capital and labor resources that can be allocated across two sectors, farming (f) and non-farming (n) production. In sector j , the total quantity of capital (K_j) or labor (L_j) can be adjusted in response to daily temperature (T_d). The productivity of that factor may change with temperature as well. We treat capital productivity (A_j^K) separately from labor productivity (A_j^L) because they may respond to temperature differently for technological reasons. In this economy, the quantity of farming output produced in a single day can be stylized as

$$q_f(T_d) = (A_f^K(T_d)K_f(T_d))^\alpha (A_f^L(T_d)L_f(T_d))^{1-\alpha}, \quad (1)$$

where both the quantity and productivity of labor and capital might respond to that day’s temperature. Similarly, we describe non-farming output for a single day as

$$q_n(T_d) = (A_n^K(T_d)K_n(T_d))^\omega (A_n^L(T_d)L_n(T_d))^{1-\omega}. \quad (2)$$

If daily temperature affects output in these sectors, then agents might respond to temperature changes by reallocating capital and labor in a way that maximizes profit. However, even if agents

reallocate resources in a sophisticated way, temperature may still constrain the production possibility frontier of the economy by altering the productivity and quantity of primitive units (capital and labor) available for production. In this analysis, we are interested in the shape of the production possibility frontier with respect to daily temperature, after agents in an advanced economy have optimized their resource allocation in (1) and (2).

On a given day, the value produced in the economy conditional on temperature is simply the combined output of the farming and non-farming sectors, each multiplied by the price p_j , which might also change with temperature.

If we could observe aggregate economic output on a daily basis, then it would be straightforward to recover the response of daily production to daily temperature. Unfortunately, because of the cost involved with collecting this data, we generally only have data on aggregate production averaged over intervals substantially longer than a day.

We are nonetheless able to recover the effect of daily temperature by making the simple assumption that the output process each day is the same. That is, we conceptualize the economy as one where each day agents observe the weather, adjust factor allocations in response, and then produce output, which is sold. Over the course of a year, this sequence occurs 365 times, and annual output is the sum of these 365 daily outputs. Indexing days by d , we can write annual earnings as the sum of earnings from each day in the year

$$annual_earnings = \sum_{d=1}^{365} [p_f(T_d) \cdot q_f^*(T_d) + p_n(T_d) \cdot q_n^*(T_d)], \quad (3)$$

where we write q^* to indicate that agents are optimizing production conditional on temperature. Written this way, it is clear that annual output depends on the full distribution of temperatures across days in the year and the extent to which temperatures on those days alter into the quantity, productivity, and allocation of primitive factors in the economy, as well as output prices.

In addition to reallocating capital and labor based on how temperature affects the productivity of these factors, it is also possible that agents can directly alter how sensitive factor productivities are to temperature by making costly defensive investments. For example, irrigation might be built to minimize the sensitivity of crops to high temperatures or air conditioners could be installed to maintain the productivity of workers. In both cases, greater defensive effort is rewarded with a reduction in the temperature-sensitivity of a factor. Following Hsiang and Narita (2012), we denote defensive adaptation effort $e \in [0, 1]$ and assume

$$\frac{\partial A_j}{\partial T} = h_j(T_d)(1 - e_j) \quad (4)$$

where $h_j(T)$ describes the sensitivity of a factor under baseline conditions that would be observed if agents exerted no adaptive effort. If agents adapt vigorously such that $e_j \rightarrow 1$, then a primitive factor becomes fully insulated from the environment and $\partial A_j / \partial T_d \rightarrow 0$. However, even in a wealthy economy this may not always happen because these defensive investments come at a cost $c(e)$ that

is increasing in adaptation effort e . For example, running an air conditioner to protect workers requires incurring electricity costs. Thus, in equilibrium, agents will select an optimal defensive effort e^* that balances the benefits of reduced temperature sensitivity of factors with the costs of the effort. If adaptation effort costs are sufficiently convex, then e_j^* will be bounded below one (Hsiang and Narita (2012)).

Finally, agents may cope with temperature-induced changes in production by sending or receiving transfers. For example, insurance might be used to offset medical costs associated with high temperature days (Deschênes, Greenstone, and Guryan (2009)). We denote net transfers in equilibrium by τ^* .

Each day, conditional on that day’s temperature, agents adapt by allocating capital and labor across sectors, exerting costly defensive effort, and sending or receiving transfers. Net income may or may not change in response to the combined effect of these actions. Augmenting Equation 3 with costs from defensive investments (benefits are implicitly captured by q_j^*) and net transfers, we see that in an advanced economy annual income is

$$annual_income = \sum_{d=1}^{365} \left[\underbrace{p_f(T_d) \cdot q_f^*(T_d)}_{\text{daily farm income}} + \underbrace{p_n(T_d) \cdot q_n^*(T_d)}_{\text{daily nonfarm income}} - \underbrace{c(e^*(T_d))}_{\text{defensive costs}} + \underbrace{\tau^*(T_d)}_{\text{transfers}} \right], \quad (5)$$

where any changes caused by daily temperature result either directly or indirectly from the structure of $h_j(T_d)$ in Equation 4, the fundamental sensitivity of the primitive factors in the economy. It is often argued that in wealthy and technologically sophisticated economies, the offsetting effects of price changes, vigorous defensive investments, and effective risk transfer mechanisms cause the net effect of temperature on income to be zero, i.e. that the derivative of Equation 5 with respect to temperature is zero. We test this hypothesis directly in the modern United States, arguably the most advanced economy ever to exist.

3 Data

We match weather, income, and transfer data at the county level for the lower 48 states during the period 1969-2011. All income and transfer measures are inflation adjusted to 2011 dollars and converted to per capita terms. Summary statistics for key variables are presented in Table 1.

Weather data To measure daily maximum and minimum temperatures as well as precipitation, we use daily surface data from the National Climatic Data Center (NCDC).² We match weather stations to counties using each station’s reported latitude and longitude. We omit observations where the maximum or minimum temperature exceeds 60 degrees Celsius or is lower than -80 degrees Celsius, as these are likely errors. If there are multiple stations within a county, we average their measures for each day. Our preferred measure of daily temperature is a simple average between

²Publicly available from ftp://ftp.ncdc.noaa.gov/pub/data/ghcn/daily/by_year/.

the maximum and minimum temperatures, which is the standard measure for average temperature during a 24-hour period.³ As discussed in Auffhammer et al. (2013), weather station data is often incomplete, sometimes due to mechanical failures, political events, or financial constraints. We drop county-by-year observations that do not have a complete set of daily weather observations. This results in a reduced sample size, with coverage that is displayed in Appendix Figure A1. Thus, our results represent the average effect of temperature on income, conditional on whatever circumstances allow counties to provide a complete record of daily weather within a single year.

Income data To measure income, we use Regional Economic Information System (REIS) data, published by the Bureau of Economic Analysis (BEA). The BEA, in turn, uses a variety of sources to construct these measures.⁴ The most inclusive income measure at the county level is *total personal income*. It encompasses all sources, including all types of labor income; proprietors' income; dividends, interest, and rent payments; and government transfer payments. A subset of personal income, *earnings*, includes only wages and salaries, other labor income, and proprietors' income. In turn, wages and salaries include tips, commissions, bonuses, and any "pay-in-kind" provided by an employer. They are measured before any deductions are taken and are derived from reports filed by employers to comply with unemployment insurance (UI) laws.⁵ Total personal income is reported on a place-of-residence basis, while wage and salary payments and other income components are reported by place of work. The residence adjustment is made using US Census estimates of worker commuting behavior. As a result, the components of personal income can sometimes exceed total personal income.

Measures of *farm income* in REIS are derived from United States Department of Agriculture (USDA) estimates, which are based on sample surveys, Agricultural Census data, and administrative data.⁶ A distinction is made between gross farm income, which includes inventory sales, and net farm income, which does not. However, additions to inventories are included in the net farm income measure. Importantly, the *net farm income* measure we use also includes transfers such as subsidies, crop insurance, and disaster payments. Our measure of gross farm income is *cash receipts from marketing crops*.

Transfers data We obtain data on total transfers from government to individuals from the REIS. These include unemployment insurance, which in turn consists primarily of standard state-

³The diurnal cycle in temperature approximately follows a sinusoid, so this standard measure is a good approximation for the true mean.

⁴For further details, see <http://www.bea.gov/regional/pdf/lapi2010.pdf>.

⁵There are only five industries that are not fully subject to these laws: agriculture, railroads, the military, private education, and religious organizations. Other data are used to infer wages and salaries in the uncovered portions of these industries. Typically, an employer will report wage and salary payments by county and by industry, resulting in very accurate county-level estimates. In a few cases, an employer will file a UI report for the whole state, rather than by county. In that case, the state total will be allocated to counties based on the industry's share in each county.

⁶For some states, estimates at the state level are allocated to counties using weights derived from the Census of Agriculture. For some commodities, Agricultural Census data are interpolated to create intercensal estimates. Because these procedures may mask some impacts of weather shocks, our estimates for the effects of temperatures on farm income should be viewed as a lower bound.

administered unemployment insurance schemes, but also includes unemployment compensation for federal employees, railroad workers, and veterans. Government transfers also include income maintenance (which includes Supplemental Security Income (SSI), family assistance, and food stamps), retirement and disability insurance benefits, public medical benefits other than Medicare, Medicare, veterans’ benefits, and federal education and training assistance. In addition, the United States has an extensive crop insurance program that has been greatly expanded over the past 30 years. Insurance plans are sold by private companies, but are heavily regulated and reinsured by the US government. We obtain annual county-level data on crop insurance indemnities for the years 1990–2011. These are publicly available from the Risk Management Agency (RMA) of the USDA. Finally, Congress has also passed numerous *ad hoc* disaster bills to give aid to farmers who suffered crop losses, regardless of whether they had insurance. County-level crop-related disaster payments for the years 1990–2010 are from USDA Farm Services Agency (FSA) administrative data, obtained through a Freedom of Information Act request.

4 Empirical approach

Estimating daily impacts from annual data

Because random variations in weather are clearly exogenous, the central empirical challenge we face is identifying the effect of daily variations in temperature on average daily income when we do not observe daily income directly. Instead, we observe annual measures of income and assume, as suggested by Equation 5, that annual income is a linear combination of many daily incomes that are similarly affected by temperature. This allows us to estimate the daily impact of temperature using annual income data.

To develop intuition for our approach, first consider a simplified example. Imagine there are only two days in a year when production occurs and each has a random temperature that is either “hot” or “cold.” Income on a hot day is \$1 and income on a cold day is \$2, but these values are not observed. Instead, we only observe total income at the end of the year (after both days) and we observe whether those days are hot or cold. After collecting data for several years, we see that if both days are hot then total income is \$2, if one is hot and the other cold then total income is \$3, and if both are cold then total income is \$4. Assuming the daily income-generating processes across years are similar, conditional on temperature, we could write these three facts as three equations containing two unknowns: the daily income on hot days and daily income on cold days.⁷ Solving the system recovers daily incomes on hot and cold days, even though they were not observed directly.

Now consider the more realistic case where daily income is a continuous and potentially nonlinear function $g(\cdot)$ of daily temperature T_d . Annual income Y is then the sum of 365 daily production

⁷If y_h is income on a hot day and y_c is income on a cold day, then the three equations are $2 = 2y_h + 0y_c$, $3 = y_h + y_c$, $4 = 0y_h + 2y_c$. Only two of these equations are needed to solve the system.

values

$$Y = \sum_{d=1}^{365} g(T_d) \quad (6)$$

where $g(T_d)$ is the function of interest equal to the bracketed term in Equation 5—the sum of earnings, defensive costs and transfers. Similar to the simpler example above, if we observe Y and the full vector of daily temperatures T_d for many years, then there will be a unique estimate for $g(T_d)$ that best describes how production over a sequence of days generates the Y values we observe at the end of each year.

To recover $g(\cdot)$, we approximate it as a stepwise function by computing the average $g(\cdot)$ over a sequence of temperature bins, indexed by m . Denote the m th temperature bin Ω^m with lower bound $\underline{\Omega}^m$ and upper bound $\overline{\Omega}^m$. The average value of $g(\cdot)$ over bin Ω^m is

$$\overline{g(\Omega^m)} = \frac{1}{\overline{\Omega}^m - \underline{\Omega}^m} \int_{\underline{\Omega}^m}^{\overline{\Omega}^m} g(T_d) dT_d.$$

Our stepwise approximation of $g(T_d)$ can be written analytically as

$$g(T_d) \approx \sum_m \overline{g(\Omega^m)} \cdot \mathbf{1}[T_d \in \Omega^m], \quad (7)$$

where the indicator variable is one if T_d is in the m th interval and zero otherwise. Substituting Equation 7 into Equation 6 and switching the order of the summations we have:

$$\begin{aligned} Y &= \sum_{d=1}^{365} g(T_d) \\ &\approx \sum_{d=1}^{365} \sum_m \overline{g(\Omega^m)} \cdot \mathbf{1}[T_d \in \Omega^m] \\ &= \sum_m \overline{g(\Omega^m)} \underbrace{\sum_{d=1}^{365} \mathbf{1}[T_d \in \Omega^m]}_{\tilde{T}^m} \\ &= \sum_m \overline{g(\Omega^m)} \cdot \tilde{T}^m, \end{aligned} \quad (8)$$

where \tilde{T}^m is equal to *the number of days during the year with temperatures that fall in bin Ω^m* . Thus, annual income Y can be approximated by a linear combination of the \tilde{T}^m variables, which we can construct easily. A regression of Y on the vector of \tilde{T}^m values must return coefficients equal to $\overline{g(\Omega^m)}$, which by Equation 7 approximate the shape of the daily output function $g(T_d)$.⁸

⁸To our knowledge, this functional form was first introduced to the econometrics literature by Deschênes and Greenstone (2011), in the context of health, although a derivation motivating this specification was not originally presented. This approach is closely related to piecewise-linear degree-days models (Schlenker and Roberts (2009)) commonly used in the agronomic literature to recover daily growth functions under similar assumptions—a derivation motivating that approach is presented in the Appendix of Hsiang (2010).

Regression specification

To estimate the economic response to daily temperatures, we estimate a version of Equation 8 that is augmented to account for autocorrelation, precipitation, lagged effects of weather, unobserved heterogeneity across counties, and nonlinear time trends. Specifically, using our panel of US counties, we estimate

$$Y_{it} = \rho Y_{i,t-1} + \sum_m \left[\beta^m \tilde{T}_{it}^m + \gamma^m \tilde{T}_{i,t-1}^m \right] + \sum_n \left[\zeta^n \tilde{P}_{it}^n + \eta^n \tilde{P}_{i,t-1}^n \right] + \mu_i + \theta_t + \epsilon_{it}, \quad (9)$$

where counties are indexed by i and years are indexed by t . Y_{it} are economic outcome variables such as *log income per capita*, the primary outcome of interest. μ_i is a set of county fixed effects that account for unobserved constant differences between counties, such as geography. θ_t is a set of year fixed effects that flexibly account for common trends, such as technological innovations or trends in climate, and year-specific shocks, such as abrupt changes in energy prices.

\tilde{T}_{it}^m is the number of days in county i and year t that have 24-hour average temperatures in the m th temperature bin, as indicated by the underbrace in Equation 8. Each interior temperature bin is 3°C wide. We define $\tilde{T}_{it}^{m=1}$ = the number of days when $T_d < -15^\circ\text{C}$, $\tilde{T}_{it}^{m=2}$ = the number of days when $T_d \in [-15, -12)^\circ\text{C}$, $\tilde{T}_{it}^{m=3}$ = the number of days when $T_d \in [-12, -9)^\circ\text{C}$, and so on. The top ($m = 17$) bin counts days with $T_d \geq 30^\circ\text{C} = 86^\circ\text{F}$. The $m = 11$ bin for $T_d \in [12, 15)^\circ\text{C} = [53.6, 59)^\circ\text{F}$ is the omitted category. The coefficients β^m are the parameters of interest, as they describe the marginal effect on Y of an additional day in the m th temperature bin, relative to a day with temperatures in the omitted category. In relation to Equation 8, the coefficient β^m estimates the component of the piecewise-constant approximation $\overline{g(\Omega^m)}$. \tilde{P}^n is defined similarly for daily precipitation across 12 bins. Each bin spans 40mm of daily precipitation, with the bottom bin corresponding to no precipitation and the top bin corresponding to precipitation $> 400\text{mm}$ in a day. Because temperatures and precipitation are, on average, serially correlated within a county, we include lagged values for all \tilde{T}^m and \tilde{P}^n variables to capture any possible direct effects that weather in the prior year might have on current output. For example, low rainfall in a prior year might reduce the quantity of groundwater available for irrigating crops in the current year.

The variable ϵ_{it} is a disturbance term that we assume may be arbitrarily correlated between counties within a state-by-year as well as within a given county over time. To account for this, we estimate standard errors that are clustered in two dimensions (Cameron, Gelbach, and Miller (2011)): within state-by-years and within counties. This approach accounts for both spatial correlation across contemporary counties within each state and autocorrelation within each county⁹.

Finally, $Y_{i,t-1}$ is a lagged dependent variable with serial correlation coefficient ρ . Including this term in the specification is important because there is substantial serial correlation in outcomes at the county level that is not accounted for by common trends. For example, the history of capital investments within a county affect production in subsequent years. It is known that one drawback

⁹See Fisher et al. (2012) for a discussion and analysis of this technique to account for spatial autocorrelation. See Hsiang (2010) for a discussion of simultaneously accounting for spatial and temporal autocorrelation.

of dynamic panel models, such as Equation 9, is that they are inconsistent when lagged dependent variables and fixed effects are estimated simultaneously by OLS (Nickell (1981)). However, this drawback is primarily a concern when panel lengths are short (e.g. ≤ 10 periods). We are not in this hazardous context, as our panel has 43 periods. We estimate that the magnitude of our potential bias is less than 5% of the magnitude of our point estimate, leaving us relatively unconcerned about this potential bias as it is far smaller than our uncertainty due to sampling error.¹⁰ We opt to utilize OLS because it conveys many advantages, allowing us to account for spatial autocorrelation, avoid using weak instruments, and estimate spatial lag models. For completeness, we show estimates our model without any lagged dependent variable and continue to obtain our main result.

Identification The average effect of daily temperature in the m th bin (β^m) is identified by Equation 9 if the exact number of days in that bin, relative to other years in the same county, are orthogonal to other potential confounders, conditional on all control variables. For example, the estimated effect of a 16°C day is identified by comparing a county to itself across years when the number of 16°C degree days was slightly different. Weather has systematic patterns in each location that are absorbed by county fixed effects. Random variations in those patterns give rise to small distortions in the distribution of daily temperature across years that we exploit for inference. As depicted in Figure 1, our estimates of each β^m are identified off of these random disturbances at each point in the temperature distribution within a location.¹¹ We follow Deschênes and Greenstone (2007) and Schlenker and Roberts (2009) in assuming that these detrended year-to-year random variations within each county are uncorrelated with year-to-year variations in other important factors that affect income.

5 Results

We first establish that daily temperature affects income in US counties and then demonstrate the robustness of this result. We next look for evidence of adaptation over time, examine the effect of weekday vs. weekend temperature, and estimate the structure of spatial lags. Finally, we compute the total effect on income accounting for dynamic impacts and separately examine different components of income to understand what mechanisms might be driving this result.

Main result

Panel A of Figure 2 presents our main result: log personal income per capita increases slightly as temperatures rise from cool to moderate, then declines approximately linearly at temperatures

¹⁰Nickell (1981) derives that the bias scales like $\frac{-(1+\rho)}{(T-1)}$, where T is the number of periods. Based on our estimate that $\hat{\rho} = 0.825$ for log personal income per capita, this number is approximately 0.045.

¹¹The effects of weather conditions over intervals longer than a day will be reflected in our estimates, but we do not identify them separately. For example, the effect of a heat wave with five 28°C days will be captured by the coefficient on the 27-30°C temperature bin, but we do not estimate a separate additional effect (eg. a “heat-wave effect”) for this specific sequence of daily temperatures.

above 15°C (59°F). Relative to a day with an average temperature of 15°C (59°F), a day at 29°C (84.2°F) lowers annual income by roughly 0.065% (−0.00065 log points). This effect is highly statistically significant.¹²

If output were uniform across 365 days in a year, then each day would contribute $\frac{1}{365} = 0.27\%$ of annual income. Thus a decline of 0.065% of annual income from a single day at 29°C (84.2°F) indicates that day is roughly 23.6% less productive than an average day. Linearizing the effect of temperature relative to the approximate zero effect at 15°C (59°F), this is a marginal change in daily productivity of $\frac{-23.6\%}{14^\circ\text{C}} = -1.68\%/^\circ\text{C} = -0.93\%/^\circ\text{F}$.

Interestingly, this *county-level* effect is quantitatively similar to the $-1.0\%/^\circ\text{C}$ and $-2.5\%/^\circ\text{C}$ decline in annual average *country-level* GDP reported in Dell, Jones, and Olken (2012) and Hsiang (2010), respectively. However, these earlier estimates were in response to average annual temperatures, not daily temperatures, and they were obtained by examining only low or middle income countries, not the United States. Despite the fact that the modern United States is the largest, richest, and most advanced economy to date, daily temperatures continue to have a quantitatively large impact on county-level income.

We next check the robustness of this result by examining lags, earnings, alternative specifications, and measurements in levels.

Temporal displacement Panel B of Figure 2 displays the estimated effect of daily temperatures on annual income per capita the *following* year. We estimate these effects jointly with the contemporaneous effect shown in panel A, as specified in Equation 9. Except for the single coefficient in the hottest temperature bin ($> 30^\circ\text{C}$), we do not observe any statistically significant effect of daily temperatures on income the following year. It is possible that the significant coefficient in the top temperature bin is spurious; because we are testing sixteen coefficients, it would not be unlikely for one to be spuriously significant. However, it is also possible that this effect is meaningful and indicates that some of the income lost from the hottest days is displaced into the following year. The estimated magnitude of this positive lagged effect is half the magnitude of the negative contemporaneous effects, indicating that roughly half of the income loss from the hottest days might be made up in the following year.¹³ There is no statistically significant evidence for temporal displacement of income for days below 30°C.

Earnings Earnings make up the majority of personal income. In panels C-D of Figure 2 we display the effect of daily temperature in current and prior years, respectively, on earnings per capita. Qualitatively, the structure of the earnings response is very similar to the income response, although the magnitudes of the point estimates are larger. Relative to a day at 15°C (59°F), a day at 29°C (84.2°F) lowers annual earnings by roughly 0.11%. Again assuming uniform output across

¹²We do not find any significant response of personal income or other income components to rainfall. Estimates are available upon request.

¹³In later simulations, we include this positive lag to ensure we do not overestimate the effect of high temperature days.

365 days, this estimate suggests that the hotter day results in roughly 40.0% lower daily earnings. This represents a linear decline of daily earnings at a rate of roughly 2.9%/°C above 15°C. Similar to total income, we see no systematic response of earnings to daily temperatures in the prior year, except possibly to very hot days with average temperatures exceeding 30°C.

Model specification In columns 1-3 of Table 2, we alter the structure of our econometric model to examine whether our main result is robust to modeling choices. In column 1 we present coefficients from our benchmark model, which includes county fixed effects, year fixed effects, precipitation, lagged weather variables, and a lagged dependent variable. In column 2, we remove the lagged dependent variable and find that the structure of the response remains unchanged, although some coefficients become somewhat smaller in magnitude (by roughly 30%) and others become somewhat larger in magnitude (by roughly 50%). In column 3, we keep the lagged dependent variable and introduce state-by-year fixed effects in place of common year fixed effects. Work by Fisher et al. (2012) suggests that this approach removes too much of the identifying variation in this context, causing attenuation bias to overwhelm results.¹⁴ Indeed, we find that the magnitude of our results is substantially attenuated (roughly 60%), but the reduction in output in both of the two hottest temperature bins remains statistically significant and the qualitative structure of the response remains the same.

Measurement in levels In column 1 of Table 3, we re-estimate our model in income levels rather than in logs.¹⁵ We find that the overall structure of the income response to temperature is unchanged and that altering a day’s temperature from 15°C (59°F) to 29°C (84.2°F) reduces annual county income per capita by \$16.71 on average. This is very similar in value to the \$17.4 dollar average loss one obtains by multiplying the 0.065% average loss (from column 1 of Table 2) by the \$26,806 average per-capita income of counties. In column 4 of Table 3, we compute population-weighted estimates that describe the effect on an average person rather than an average county. These effects are smaller and are not statistically significant in the top temperature bins, although warm days (24-27°C or 75.2-80.6°F) are estimated to cost the average individual \$4.80. This change in the response function reflects a re-weighting of the relative contributions of farm

¹⁴Temperature is strongly correlated across moderate distances, so when a specific county has warm temperatures it is likely that neighboring counties also have warm temperatures (this spatial correlation motivates our use of standard errors that are clustered by state-year). The inclusion of state-year fixed effects thus absorbs all county-level temperature variation that is spatially correlated within each state-year, variation that is “real” in the sense that it is very unlikely to be driven by measurement errors at individual stations. However, inclusion of state-year fixed effects does not generally reduce the quantity of measurement error at the county level because these errors are likely independent across weather stations and counties. Thus, inclusion of state-year fixed effects substantially reduces the real partial variance in temperature (σ_T^2) that is described by our independent variables without reducing the variance in measurement error (σ_e^2). Since attenuation bias causes coefficient estimates to scale by $\frac{\sigma_T^2}{\sigma_T^2 + \sigma_e^2}$, parameter estimates are driven towards zero as the model becomes increasingly saturated because $\sigma_T^2 \rightarrow 0$ while σ_e^2 remains relatively unchanged. See Fisher et al. (2012) for a complete treatment of this issue. Those authors conclude that models similar to ours are over-saturated when state-by-year effects are included.

¹⁵We Winsorize incomes at the 1st and 99th percentile to prevent outliers from exerting strong influence on our estimates.

and non-farm income to the estimated effect, discussed in greater detail below.

Having established that the adverse effect of warm and hot temperatures on income is robust to modeling choices, we next consider whether there has been any adaptation during our time period, how the effect of temperatures differs on weekends versus weekdays, whether there are any spatial spillovers, and how income dynamics affect the total impact of temperature.

Extensions of the main result

Adaptation and stability over time We are interested in whether daily temperature affects economic outcomes in the modern United States. However, it is possible that our result is driven by earlier observations in our sample, which stretches back to 1969, and that the behavior of the most recent years is not well-described by the pooled estimate. For example, residential air conditioning spread steadily through the United States between 1960-2000, a fact that is likely responsible for reducing heat-related mortality during the last century (Barreca et al. (2013)). It is possible that this or other adaptations have altered the sensitivity of income to daily temperatures in a substantive way. To examine this issue, we estimate the response of income to temperature for each decade separately, shown in Figure 3. These estimates are noisier because each relies on a much smaller sample, but they do not differ substantively from our pooled estimate. This suggests that adaptations during the last half-century have not fundamentally altered the relationship between temperature and income in the United States and that our main result remains valid during the first decade of the twenty-first century.

Weekdays vs. weekends Most economic activity in the United States occurs on weekdays rather than weekends. Some sectors, such as entertainment or retail, may generate greater value on weekends, but many of the sectors where high temperatures are likely to have adverse impacts focus production on weekdays. Thus, it is a useful exercise to examine whether hot weekdays or hot weekends drive our result. To do this, we estimate a variant on Equation 9 where we jointly estimate the marginal effects on annual income per capita of an additional weekday and an additional weekend day in each temperature bin. We plot the response of income to weekday temperatures and weekend temperatures in Figure 4 (coefficients are reported in columns 4a-b of Table 2).

We observe that the negative effect that high temperature days have on annual income is driven by hot weekdays and not hot weekends. Warm weekend days with temperatures up to 27°C (80.6°F) have very small effects with narrow confidence intervals, allowing us to reject the hypothesis that they have large effects on income. Hot weekend days with temperatures above 27°C (80.6°F) have a positive and significant effect on income, increasing annual income by roughly 0.05% on the hottest days. It is possible that hot weekends benefit annual income because hot temperatures compliment some of the industries where production is focused on weekends. For example, consumers may shop in air-conditioned malls or eat in outdoor restaurants more on hot summer days than at other

times of the year. Consistent with this notion, very cold weekends with temperatures below -3°C (26.6°F) lower annual income.

In contrast, hot weekdays reduce annual income by 0.087% per day above 27°C (80.6°F), an effect substantially larger in magnitude than the 0.065% effect from the estimate that pooled weekends and weekdays. Under the conservative assumption that all economic value is produced on weekdays, this indicates that these hot weekdays are roughly 22.7% less productive than average weekdays at the optimal temperature. If we assume that economic value is produced uniformly across days in the week, including weekends, then these hot weekdays are roughly 31.8% less productive than average weekdays at the optimum temperature. For days above 30°C (86°F), analogous productivity losses are 32.7% and 45.8%, respectively.

Spatial lags Similar to how we use a temporal lag structure to look for temporal displacement, we look for spatial displacement using a spatial lag structure. We assume that a county’s income is affected by both the number of days that it is exposed to a specific temperature and the average exposure of neighbors whose centroids fall within concentric annuli (around the county’s centroid) with 100km widths out to a maximum distance of 500km (Cressie and Wikle (2011)). If high temperature days cause output to be displaced to neighboring counties, then the coefficients on neighboring temperatures will be positive. Figure 5 displays the spatial lag structure for 24- 27°C days, 27- 30°C days, and $> 30^{\circ}\text{C}$ days (jointly estimated), indicating how county income is correlated with its own temperature and its neighbors’ temperature at various distances. We find that high temperatures continue to have a negative effect on own income even when accounting for neighbors’ temperature and that there are no significant positive benefits of being near hot counties. If anything, there is some evidence that high temperatures in neighboring counties have negative effects on a county’s own income. This suggests that either there are negative spillovers across counties that experience high temperature days or that neighbors’ temperatures are a proxy measure for some other temporary environmental condition that negatively affects income but is not captured by our benchmark model, such as the length of hot spells.

Accounting for dynamic effects on income Annual income in US counties is serially correlated, so a temperature-driven income change in year t will have an indirect effect on income in year $t + 1$. In Equation 9, we estimate that $\hat{\rho} = 0.825$, implying that $Y_{i,t+1} = 0.825Y_{it} + X_{it}\beta + \epsilon_{it}$. Thus, a temperature-driven income loss of \$1 in year t will result in an income loss of \$0.825 in year $t + 1$, \$0.68 in year $t + 2$ and so on, relative to a counterfactual income trajectory where no loss was suffered in year t . The discounted sum of these income losses between the time of the temperature event and ∞ represents the net present value (NPV) of lost income attributable to the temperature event. By computing the NPV of the difference in county i ’s income (ΔY_{is}^m) at each moment s that was attributable to the temperature event $\tilde{T}_{it}^m = 1$ at time t , we can estimate

the full net present value of a day with $\tilde{T}_{it}^m = 1$. Using a discount factor δ , this is

$$NPV(\text{total_lost_income}) = \sum_{s=t}^{\infty} \delta^{(s-t)} \Delta Y_{is}^m \approx \sum_{s=t}^{\infty} \delta^{(s-t)} \hat{\rho}^{(s-t)} \hat{\beta}^m = \frac{1}{1 - \hat{\rho}\delta} \hat{\beta}^m. \quad (10)$$

Thus, the NPV of the altered income trajectory is a linear scaling of coefficients by $\frac{1}{1 - \hat{\rho}\delta}$. Using a discount factor $\delta = 0.95$ (implying an annual discount rate of 5%) and $\hat{\rho} = 0.825$, we estimate this scaling factor to be 4.62. In NPV terms, this implies that a warm day at 22.2°C (72°F) costs $4.62 \times -0.000294 = -0.00136$ log points of annual per capita income (relative to a 13°C (55.4°F) day) because, in addition to altering contemporaneous income, it alters a county’s future income trajectory. A hot day exceeding 30°C (86°F) is estimated to cost $4.62 \times -0.000757 = -0.0035$ log points of annual per capita income in NPV terms.¹⁶ Recalling that a randomly selected day is responsible for $\frac{1}{365} = 0.00274$ log points of annual income, the NPV of the total cost of a warm or hot day is roughly 0.5 and 1.3 days’ worth of average income, respectively.

We next consider the different components of income to try and understand the mechanisms that might contribute to this result.

Components of income

Much research on the economic impact of temperature in the US has focused on farming, since the negative impact of adverse weather on crop yields is acute (Mendelsohn, Nordhaus, and Shaw (1994), Schlenker, Hanemann, and Fisher (2005), Deschênes and Greenstone (2007), Schlenker and Roberts (2009), Welch et al. (2010), Fisher et al. (2012), Burke and Emerick (2013)). However, recent work has indicated that non-farm income sources outside of the US may also be affected by high temperatures (Jones and Olken (2010), Hsiang (2010), Dell, Jones, and Olken (2012)). These authors suggest that this non-farm effect may be driven by the well-documented productivity decreases of workers who are exposed to thermal stress (Mackworth (1946), Froom et al. (1993), Seppanen, Fisk, and Lei (2006)).

In the US, high daily temperatures are associated with reductions in yields of major crops (Schlenker and Roberts (2009)) as well as reductions in labor supplied among workers exposed to outdoor temperatures, which includes manufacturing (Graff Zivin and Neidell (2014)). These studies demonstrate that the productivity of crops (A_f^K in Equation 1) and the quantity of labor supplied (L_n in Equation 2) depend on daily temperature and suggest mechanisms that might explain our main finding that high temperatures reduce income in US counties. However, without observing price changes, these studies alone are not conclusive. If prices change with temperature, as is possible in Equation 3, then changes in production might not translate into changes in earnings. To consider whether these mechanisms might be contributing to the effect we document

¹⁶For a discount rate of 3%, a warm day at 22.2°C (72°F) costs 0.00147 log points of current income in NPV, while a hot day above 30°C (86°F) costs 0.00379 log points. At a 7% discount rate, analogous values are 0.00126 and 0.00325 log points.

above, we directly examine how the agricultural and non-agricultural components of income respond to temperature. We then compare both the structure and magnitude of these responses to earlier results by Schlenker and Roberts (2009) and Graff Zivin and Neidell (2014). To facilitate comparison, we reproduce the main results of both studies in Figures 6A and 7A, respectively.

It is worth noting here that Graff Zivin and Neidell (2014) obtain data on the quantity of labor supplied (L_n in Equation 2) but cannot observe labor effort, i.e. the productivity of labor supplied (A_n^L in Equation 2). Lab studies indicate that the labor productivity response to temperature is qualitatively similar in structure to the response reported in Graff Zivin and Neidell (2014) (Mackworth (1946), Froom et al. (1993), Seppanen, Fisk, and Lei (2006)). Thus, the total labor effects on income might be larger than the estimates in Graff Zivin and Neidell (2014) suggest, but the overall structure of the response should be similar.

Agricultural income To examine how crop losses contribute to our main result, we repeat our analysis, replacing the dependent variable with *log revenue from crop sales per capita*. In panel B of Figure 6, we plot the effect of hot days on annual income from crops and observe steep declines when daily average temperatures rise above 27°C (tabulated in Appendix Table A1). This structure is very similar to the yield response obtained by Schlenker and Roberts (2009) (Figure 6A). The slightly higher breakpoint of 29-32°C in that study and its steeper decline is likely because Schlenker and Roberts (2009) use hourly temperature, whereas our analysis uses daily averages. Because days with 24-hour average temperatures of 27°C are likely to have some hours above 29°C, we would expect to observe declines on days with average temperatures of 27°C in our analysis, even if crop yields do not deteriorate until the hourly temperature reaches 29°C. Thus, we interpret our results in Figure 6B as consistent with the crop yield response in Figure 6A reported by Schlenker and Roberts (2009). Our results suggest that higher crop prices do not dramatically offset yield losses caused by high temperature days; thus, reductions in yields translate into reductions in income.

Quantitatively, the decline in crop income explains a significant share of our main result for total income: a 30°C day reduces annual crop income by 0.523% but lowers total income by only 0.076%. This large decline in crop income is broadly consistent with the magnitude of changes reported by Schlenker and Roberts (2009), although a direct comparison is difficult because of the difference in measurement described above. The outcome in Schlenker and Roberts (2009) is the yield effect of 24 hours at each exact temperature. Because 24 hours at 35°C reduces annual yields by roughly 0.03 log points (an approximate average across the three crops in Schlenker and Roberts (2009)), one hour at 35°C should reduce annual yields by roughly $\frac{0.03}{24} = 0.00125$ log points. A day with average temperature of 29°C might have roughly one hour at this higher temperature during the day’s peak temperature, and we estimate that such a day would cause crop income to decline by 0.00187 log points. Thus, while we cannot make a perfect comparison between these two sets of results, this back-of-the-envelope calculation does seem consistent with the hypothesis that high-temperature yield declines cause a decline in income that is not offset by rising prices.

In panel C of Figure 6, we examine how *net farm income per capita* (in levels) responds to

daily temperature and find that it declines by \$21.07 for each day above 30°C .¹⁷ The structure of the response of total farm income differs somewhat from that of crop income: we observe lower farm incomes starting at temperatures around 20°C . We lack the data to determine precisely the mechanism that could mediate effects at these lower temperatures. However, the observed structure is broadly consistent with the labor productivity response discussed below. Alternatively, farmers may be increasing expenditure on inputs to combat the negative impacts of temperatures on yields. Finally, it is also possible that other productive factors respond negatively to these lower temperatures.¹⁸

Non-agricultural income We examine whether non-agricultural income might be playing a role by repeating our analysis on *log non-farm income per capita*, shown in Figure 7B (tabulated in Appendix Table A1). We see that non-farm income is relatively flat (albeit noisy) at low temperatures and then begins to decline systematically at temperatures above 15°C, the same breakpoint observed for total income (Figure 2). However, the magnitude of the effect on non-farm income is smaller, with temperatures at 25°C lowering annual non-farm incomes by only 0.021% relative to 15°C whereas the analogous loss of annual total income is 0.059%. Both of these features of the response, the smaller magnitude and the lower breakpoint temperature, are broadly consistent with the response of labor supply documented by Graff Zivin and Neidell (2014) (Figure 7A) and labor productivity responses from lab experiments (Seppanen, Fisk, and Lei (2006)). As with the crop yield response, the breakpoint documented by Graff Zivin and Neidell (2014) (~25°C) is a higher temperature than what we observe in non-farm income (15°C). This difference is likely due in part to Graff Zivin and Neidell (2014) using daily maximum temperature rather than daily average temperature as we do—although the 10°C difference might be too large relative to normal diurnal temperature variations to be fully explained by this fact alone.¹⁹ It is possible that changes in the quality of labor, i.e. the intensive margin, are responsible for this lower turning point: lab studies summarized in Seppanen, Fisk, and Lei (2006) indicate that productivity begins to decline at slightly lower temperatures (~21-22°C). We observe that the point estimate for non-farm income increases in the hottest temperature bin. However, this point estimate is noisy and is neither statistically different from zero nor from the negative estimates at the adjacent temperature bin.

Quantitatively, our estimated effect of temperature on non-farm income is roughly four times what one might expect based only on previous labor supply results, which is consistent with the notion that unmeasured labor productivity effects are comparable or larger in magnitude to documented labor supply effects. For a day with an average temperature of 25°C, annual non-farm income is estimated to fall by 0.000213 log points, which corresponds to a loss of 7.8% of an average

¹⁷Net farm income is not amendable to a log model because many observations in the sample are negative.

¹⁸We observe cash receipts from livestock sales and find that they are not significantly affected by temperature. Key, Sneeringer, and Marquardt (2014) find that, while dairy production is negatively correlated with local average temperatures, there is no relationship between temperature deviations and dairy production.

¹⁹The average difference between the daily average and maximum temperatures in our sample is about 6.5°C. A difference of 10°C is slightly above the 90th percentile in that distribution.

day’s non-farm output ($\frac{0.000213}{1/365} = 0.078$) relative to the optimum temperature. Maximum temperatures on such a day might reach low 30’s or even 35°C. Based on results reported by Graff Zivin and Neidell (2014), daily maximum temperatures in this range might result in a roughly 30-minute drop in labor supply, or 6.5% of the average 7.67 hour workday among workers who spending a significant amount of time working outdoors. Because these thermally-vulnerable workers — termed “high risk” in Graff Zivin and Neidell (2014) — constitute 28% of the national workforce (Houser et al. (forthcoming)), a randomly selected worker would on average supply 1.8% less work on this hot day, which is roughly one fourth of the 7.8% loss of non-farm income that we document.

Transfers Finally, we examine the effect of daily temperature on several types of transfers that we are able to observe, as these may change with temperature and potentially offset lost earnings (recall Equation 5). For example, risk-sharing within the US causes federal government transfers to increase following natural disasters (Healy and Malhotra (2009), Deryugina (2014)). In Panel A of Figure 8, we show the effect of daily temperatures on total annual transfers from the government (excluding crop-related payments). We estimate that daily temperatures have zero effect on total transfers. Our confidence intervals on these estimates are narrow, allowing us to reject fractional changes in transfers that are similar in magnitude to the estimated effect on income. One component of these government transfers that might conceivably respond to daily temperatures is *total spending on public medical benefits*. Daily temperatures are known to affect health outcomes (Deschênes, Greenstone, and Guryan (2009); Deschênes and Greenstone (2011); Barreca et al. (2013)), possibly increasing demand for medical care and thus for these transfers. We estimate the effect of daily temperatures on medical benefit transfers in Panel B of Figure 8 and find no statistically significant effect or pattern, although these estimates are notably noisier than the measure of total transfers.

In Panels C and D of Figure 8, we display the effect of daily temperature on *ad hoc crop disaster payments* and *crop insurance indemnities*, respectively. These data are only available for 1990-2010, and we estimate these effects in levels because many of the observations are zeroes. We observe no systematic pattern for disaster payments, although there is suggestive evidence of small reductions in payments for cold days and a small increase in payments in the hottest days ($> 30^\circ\text{C}$). Crop insurance payouts appear to offset farm income losses somewhat (recall Figure 6C), increasing steeply for days that exceed 27°C (80.6°F). For each day above 30°C (86°F), counties receive an average of \$5 per capita in crop insurance payouts. Because crop insurance payments are already included in farm income, these estimates suggest that farm income losses would be roughly 25% higher if crop insurance were not available.

6 Counterfactual simulations

To develop a sense for the overall influence of daily temperature on economic outcomes, we conduct two thought experiments where we apply our estimates from Equation 9 to counterfactual temperature distributions (shown in Appendix Figure A2). In the first, we imagine that daily temperature

is a choice parameter that counties can optimize every morning before production occurs. By comparing outcomes from this simulation with actual economic performance, we roughly characterize the economic cost that counties bear because their temperature is determined by geography rather than endogenously. In the second counterfactual, we compare economic performance when counties are subject to their historical distribution of daily temperatures (1969-1990) in comparison to the distribution of daily temperatures they are forecast to experience at the end of the twenty-first century under “business as usual” climate change (2080-2099). This comparison provides an estimate for the economic cost of temperature changes caused by climate change under the assumption that future patterns of adaptation are similar to those we observe historically.

Both simulations apply our partial equilibrium estimates for the income response to marginal changes in the distribution of daily temperatures to counterfactual settings in which changes to the temperature distribution are not marginal. Thus, our simulation results should be interpreted with caution, as there may exist general equilibrium responses that would emerge—and which our model does not capture—should the United States actually be subject to one of these counterfactuals. Despite this shortcoming, we think these simulations are a useful exercise for developing a sense of scale for the economic impact of daily temperature historically and in the future.

Temperature as a component of geography

We imagine a world in which every day, counties can choose their temperature before economic activity commences. If, just like many other inputs to production, counties optimize their daily temperature to maximize income,²⁰ they will set their temperature to 12-15°C (53.6-59°F) each day.²¹ We interpret this scenario as relaxing one of the constraints that is imposed on counties by their geography. We think of the economic cost of this geographic constraint as the difference between outcomes in this scenario and an otherwise identical scenario where counties experience the actual sequence of daily temperatures.

We first estimate our benchmark model (Equation 9) with the full sample and all variables. We use Winsorized personal income per capita as the outcome variable, weighting by population to enable us to accurately aggregate the estimates into a national income measure. Then, holding all coefficient estimates fixed and taking output in 1969 as given, we predict what economic output would have been during 1970-2011 had each county been able to fix their daily temperature at 12-15°C (53.6-59°F).²² To ensure this counterfactual is comparable to the realized output trajectory, we also use the same procedure to predict output using actual weather observations (“historic”

²⁰If counties could actually choose their daily temperature, they would likely optimize it over a combination of income effects as well as its amenity value. For simplicity, we consider only income effects here.

²¹In this thought experiment, we avoid using the corner solution at $< 5^{\circ}\text{F}$ ($< -15^{\circ}\text{C}$) as the optimum. The estimate at this edge is highly uncertain and there is substantial likelihood that the true effect at these temperatures is below the interior optimum.

²²This procedure is slightly more complicated than using the `predict` command in Stata because of dynamic effects. Output in each year depends on output in the prior year, which is affected by counterfactual temperatures. Thus, we predict each year in the sample sequentially, allowing us to condition on prior year counterfactual temperatures in each estimate.

simulation). To characterize the uncertainty of these simulations, we bootstrap 1,000 estimates of Equation 9 and simulate both temperature scenarios for each estimate. We then multiply county-level predictions by the county’s population and aggregate them to a single annual measure of total income.

Figure 9 displays the model estimates and 95% confidence intervals for simulated trajectories of aggregate United States income in both scenarios, as well as the full distribution of implied annual growth rates for each realization. We note that the historic simulation closely mimics the observed national income trajectory, so much so that they are not visually distinguishable for most of the sample. We estimate that total personal income in the United States would have grown dramatically faster had counties’ temperatures not been constrained by geography. Our model estimate is that United States income would grow on average 4.05% per year (95% confidence interval: 3.12-4.65%) if counties could have optimized their temperature to maximize income, compared to income growth of 2.36% per year (95% confidence interval: 2.20-2.39%) in the historic simulation. Actual growth during this period was also 2.35% per year. It takes one decade for the confidence intervals of the two simulations to separate completely. After 42 years, at the end of our sample period, total income is nearly 100% larger in the temperature-optimized simulation relative to the realistic simulation. Overall, these results provide *prima facie* evidence that counties’ inability to control their daily temperature costs the United States roughly 1.69 percentage points of growth per year.

We find the results presented in Figure 9 to be stark, particularly because our estimates suggest an extremely hot day has a relatively modest cost of about \$20 per capita per day. Thus, reconciling these seemingly different magnitudes merits discussion. There are two reasons why income trajectories diverge so rapidly from historical experience when we simulate counties optimizing their temperature. The first is that temperatures are suboptimal on most days for most counties; average temperatures fall in the 12-15°C range 9% of the time. Thus, even if the marginal per capita loss from an additional warm or cool day is only a few dollars, hundreds of such days across thousands of counties per year add up to substantial losses nationally. This is analogous to the proverbial “death by a thousand cuts”—small economic costs imposed by geography accumulate to large quantities because economies are continuously exposed to these losses. Hsiang and Jina (2014) describe a similar phenomenon when examining the effect of repeated tropical cyclones on economic growth. The second reason why small daily effects of temperature accumulate to large aggregate growth effects is because the income path of counties has a dynamic component that depends on income in prior years, so economic losses imposed in an early period grow over time.

It is impossible to test whether these results would hold in reality, and it is likely that general equilibrium adjustments would alter these impacts. We do not interpret these values too literally and stress that they are most useful as a thought experiment that provides an order-of-magnitude estimate. Overall, we think the general structure of these results force us to carefully consider the potentially large importance of this single environmental variable for the world’s richest economy.

Climate change

To consider the potential economic impact of climate change using our main result, we simulate economic production under two scenarios. In the first scenario, we compute the average within-year distribution of daily temperatures and daily rainfall for each county during 1969-1990, the pre-warming period. We then predict what economic production would be assuming that each county is exposed to exactly this average weather distribution each year over a 20-year period. Because we want to abstract from specific events during the time period of our data, we ignore the estimated year fixed effects in this simulation.

In the second scenario, we replace the average temperatures during 1969-1990 with projections of the average within-year distribution of daily temperatures in each county during 2080-2099 under the “business as usual” warming scenario known as the Representative Concentration Pathway (RCP) 8.5. To account for uncertainty in these projections of future warming (Burke et al. (forthcoming)), we use 44 model projections developed by Houser et al. (forthcoming) that provide a range of estimates for total warming levels and spatial patterns of warming conditional on a uniform greenhouse gas emissions scenario (RCP 8.5). Differences between these climate model projections emerge from numerous differences in modeling decisions and techniques used by climate modeling teams around the world, with the set of model outputs representing the collective research effort of tens of thousands of researchers across multiple decades. By employing all 44 projections, we are able to construct 44 possible county-level distributions of daily temperature under this warming scenario, each of which generates a unique income trajectory for that county. We aggregate county-level outcomes within each projection to a national income trajectory that accounts for potential spatial covariances in each warming scenario.

To compute the effect of climate change on the behavior of national income growth, we compare the 20-year trajectory of aggregate incomes in each of these 44 warming scenarios against the 20-year trajectory simulated using each county’s 1969-1990 climatology. These 20-year simulations provide perspective on the potential impact that warming may have on economic growth, but they do not provide estimates for the total integrated impact of climate change over the 21st century. Estimating the integrated impact of all warming experienced between 1990-2100 would require county-level projections of daily events for each model over the entire 111 year period.

Figure 10A displays the estimated effect of warming over 20 years on national average income relative to a pre-warming scenario. Each line represents the estimated effect of warming using one of the 44 model simulations. After being exposed to an altered distribution of daily temperatures for 20 years, our estimates suggest that average incomes would be lower by 2.31% relative to the counterfactual with no warming from the 1969-1990 baseline. The full range of projections spans -3.12% to -1.17% and the inner 36 projections (90% of simulations) span -2.90% to -1.72%.

We then estimate the distribution of effects that warming is expected to have on the growth rate of income. For each of the 44 model simulations, we compute the rate of per capita income growth and difference it from per capita income growth in the baseline scenario. We plot the distribution of these differences in Figure 10B. Following the approach used by Houser et al. (forthcoming), each of

the 44 estimates has been up-weighted or down-weighted so that the collection of climate scenarios mimics the distribution of global warming outcomes that we expect based on estimates for the sensitivity of the global climate to greenhouse gas emissions. Importantly, this procedure captures some of the impact of the “fat tail” in the global climate sensitivity discussed in Weitzman (2009). Thus, the distribution of growth rate effects shown in Figure 10B can be literally interpreted as the probability distribution of impacts that we expect (based on our reduced-form estimates) if emissions follow the “business as usual” scenario described in RCP 8.5. Our median estimate is that growth will slow by 0.12 percentage points per year under this warming scenario, with the central 50% of probability mass (the inter-quartile range) spanning -0.13 to -0.11 percentage points, the inner 90% of mass spanning -0.15 to -0.09 percentage points and the full range of projections spanning -0.16 to -0.06 percentage points.

Similar to the temperature-optimization simulations above, these partial-equilibrium estimates for the income growth effect of warming do not account for any general equilibrium adjustments. Furthermore, these estimates do not account for any novel types of adaptation based on innovations that were not available to counties in our sample. However, these estimates also do not account for other changes in counties’ climates, such as changes in rainfall, natural disasters, or ecology. Nor do they capture climate-related economic events in foreign countries that affect domestic incomes. Thus, there are reasons to believe that the actual economic impact of climate change on the United States may be smaller than our estimates suggest as well as reasons to believe the impact will be larger.

It is useful to compare our estimates to results from Houser et al. (forthcoming), who take a different but complementary approach of estimating the growth effect of climate change. Houser et al. use sector-specific, county-level dose-response functions for changes caused by daily climatic conditions to estimate economic impacts. For example, they compute the expected hours of labor supply that are lost due to high temperatures in both high- and low-risk industries. These impacts are then incorporated into a general equilibrium model of the US economy and changes in the growth rate are computed under the same 44 climate scenarios that we use here. 90% of the probability mass for growth changes in Houser et al. falls between -0.016 and -0.094 percentage points per year (median: -0.042). We compare these to our estimates in Figure 10B. The two distributions overlap substantially and are of a similar magnitude, suggesting that our reduced-form approach – which does not model sectors specifically – obtains results broadly consistent with an approach that does. The fact that our estimates tend to be larger than those of Houser et al. is likely because the latter omits many factors that are known to be affected by the climate, e.g. labor productivity, but which did not meet certain criteria for inclusion in that study.²³ Our reduced-form estimates account for all pathways through which daily temperature affects income; thus, our estimates capture certain impacts that were omitted from Houser et al.

Our estimated effect of future warming on growth is much smaller than our estimated effect for

²³For example, mechanisms were required to have strong evidence from the econometric literature before they were included in the analysis of Houser et al. (forthcoming).

the historical contribution of temperature as a component of geography. There are three reasons for this. The first is mechanical: we do not extrapolate our response function beyond the support of historical daily temperatures. However, under the climate change scenarios, many days will exceed historical maximums. Sticking with the original definition of our temperature bins causes the marginal income effect of all days exceeding 30°C (86°F) to be capped at -0.0757 log points. If the response function pictured in Figure 2A continues to slope downward at very high temperatures, then our estimates for future impacts will be too small. The second reason that climate change has a smaller effect than historical geography is simply that the environmental change that occurs in our temperature-optimizing thought experiment is larger, from an economic standpoint, than the environmental change under global warming. Most daily temperatures are fairly far from the optimal temperature: the average distance from the optimum temperature is 9.01°C (16.2°F). By contrast, the average increase in the county-level temperatures under our climate change scenarios is only 5.2°C (9.36°F). Thus, allowing counties to remove all the dispersion in their daily temperatures is more dramatic than shifting the distribution of temperatures to warmer temperatures due to climate change (see Appendix Figure A2). The third reason is that some counties benefit from warming in the climate change scenario, since cold temperatures also appear to suppress productivity, although not as much as warm temperatures. Reducing the number of county-days that are initially below the optimum temperature increases those counties' output, an effect that partially offsets the losses from warming up the days that are initially optimal, warm, or hot. In contrast, all 91% of county-days that are at initially suboptimal temperatures benefit from having their temperatures shifted to the optimum and there are no offsetting effects to attenuate the difference between the historical and optimal simulations.

7 Discussion

Forward-looking and well-resourced producers will adapt to their environment through factor reallocation and defensive investments, while price adjustments may partially offset quantity shortfalls. Furthermore, individuals may insure income via transfers. However, all of these adjustments have costs, so optimal adaptation will equalize the marginal cost of adjustment with the marginal benefit of reducing environmentally-caused losses. If environmental changes reduce income in equilibrium, the implication is that the marginal costs of adaptation are sufficiently large such that it is optimal for economic production to not be fully insulated from the environment. We find that even after agents fully adjust, daily temperature affects the production possibility frontier of counties in the United States, suggesting that even in a wealthy and sophisticated economy, the costs of adaptation prevent it from being fully utilized.

Contrary to the notion that wealthy countries easily decouple their economy from the environment because they “have the resources to adapt,” the framework we present here suggests that the quantity of optimal adaptation is determined by the marginal costs and benefits of adaptation technology, not simply by the wealth of the population. Some prior evidence supports the

notion that adaptation effort with respect to climatic conditions should not depend on income: Anttila-Hughes and Hsiang (2011) and Hsiang and Jina (2014) observe that income losses following hurricanes do not vary by income within and across countries, respectively. However, total wealth could be important for allowing adaptation to occur if agents are credit constrained and existing assets are an important resource to draw on for defensive investments or insurance. Thus, while our results suggest that wealth is not a sufficient condition to entirely decouple an economy from its local environment, it remains likely that poorer economies are *more* affected by their environment than wealthy economies. Do credit constraints explain the large differential effect of temperature on rich and poor economies characterized by Dell, Jones, and Olken (2012)? It is certainly possible, and this question is an important area for future research. However, we note that in addition to credit constraints, other explanations may also be playing a role in the apparent differences between environmental sensitivities of rich and poor countries.

One hypothesis is that countries being poor and having their income tightly coupled to the environment is correlated but not causal. This could occur if the same factors that cause individuals to be poor also make their incomes sensitive to the environment. For example, economies dominated by agriculture might generate low incomes because agriculture is less profitable than manufacturing or services. These agricultural incomes may also be more strongly influenced by environmental conditions, as has been documented in many contexts, including ours. One interpretation of this observation is that it is simply more difficult to adapt agricultural production to changing climatic conditions than it is to adapt manufacturing or services. For example, Hsiang and Narita (2012) argue that the high spatial density of human and physical capital used in manufacturing and services makes it more cost effective to invest in defensive infrastructure in these sectors, relative to the low spatial density of capital in agriculture. If it is the high marginal cost of adaptation in agriculture that makes poor countries more vulnerable to the climate, rather than poverty itself, this theory would suggest that we may increase the climate resilience of poor countries with policies that bring down the cost of adaptations in agriculture. In a second example where poverty and climate vulnerability are correlated but not causally related, countries might be poor because they have weak property rights regimes and unfavorable political institutions. These factors might also make them relatively more susceptible to environmental conditions for a variety of reasons (Kahn (2005)), including, for example, by altering the incentives to invest in defensive infrastructure.

A second hypothesis is that reverse causality is playing a role in generating observed patterns, i.e. some populations are poor precisely because they are more susceptible or more heavily exposed to environmental changes in ways that are not captured by average temperature. Many poor countries are in tropical and subtropical climates where marginal changes in temperature might have unique economic implications. For example, the generally negative correlation between rainfall and temperature in these regions (which makes heat more damaging for agriculture) reverses sign at higher latitudes (Auffhammer et al. (2013)), and tropical temperature changes are dominated by the El Niño-Southern Oscillation, which has global economic and political impacts (Hsiang, Meng, and Cane (2011)). In this analysis, we observe that the effect of temperature on income

is nonlinear, with more negative marginal effects appearing at higher temperatures. If a similar structure occurs in other contexts, it is plausible that countries with higher baseline temperature will suffer larger economic penalties on average because their daily temperature distributions are centered at higher temperatures.²⁴

Our counterfactual simulation where counties optimize their temperature suggests that the current climate of the United States imposes a substantial economic cost, although there is not an obvious way in which these costs can be avoided. Because there are no comparable prior estimates for this cost, it is understandably difficult to gauge the reliability of these estimates. The only loosely comparable estimate is from Hsiang and Jina (2014), who estimate the global economic burden imposed by hurricanes using a similar technique and find similarly large costs.²⁵ While it may be tempting to dismiss both estimates as implausible, we emphasize that these findings are indeed reasonable when one considers the pervasiveness of suboptimal climatic conditions. In this analysis, counties are exposed to suboptimal temperatures 91% of days. Because income accumulates and compounds with time, the total loss grows to be substantial.

While extraordinary achievements in science, technology, politics, and social institutions over the last millennia have lifted modern economies to levels never before achieved, we find that these forces are constantly in opposition to at least one environmental factor, temperature, that continuously slows down economic progress. The data suggest that this “environmental friction”—which we estimate will worsen under climate change—slows income growth because many of the building blocks of the economy are each regularly exposed to suboptimal environmental conditions and completely adapting to these conditions remains costly at the margin.

²⁴Dell, Jones, and Olken (2012) investigate this idea using annual temperatures and a dummy variable for “hot” countries and find limited evidence to support this notion, although it is possible that more granular analysis similar to this study may reveal nonlinearities.

²⁵Estimates by Hsiang and Jina (2014) suggest global GDP growth is reduced roughly 1.3% per year by the cumulative income effect of thousands of tropical cyclones.

References

- Acemoglu, Daron, Simon Johnson, and James A. Robinson. 2002. "Reversal Of Fortune: Geography And Institutions In The Making Of The Modern World Income Distribution." *The Quarterly Journal of Economics* 117 (4):1231–1294.
- Anttila-Hughes, Jesse K. and Solomon M. Hsiang. 2011. "Destruction, Disinvestment, and Death: Economic and Human Losses Following Environmental Disaster." *Working paper* URL http://conference.nber.org/confer//2012/EEEHC12/Hsiang_Antilla-Hughes.pdf.
- Auffhammer, Maximilian, Solomon Hsiang, Wolfram Schlenker, and Adam Sobel. 2013. "Using weather data and climate model output in economic analyses of climate change." *Review of Environmental Economics and Policy* 7 (2):181–198.
- Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S Shapiro. 2013. "Adapting to climate change: the remarkable decline in the US temperature-mortality relationship over the 20th century." Tech. rep., National Bureau of Economic Research.
- Burke, Marshall, John Dykema, David B Lobell, Edward Miguel, and Shanker Satyanath. forthcoming. "Incorporating Climate Uncertainty into Estimates of Climate Change Impacts." *Review of Economics and Statistics* .
- Burke, Marshall and Kyle Emerick. 2013. "Adaptation to climate change: Evidence from US agriculture." *Working paper* .
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller. 2011. "Robust inference with multiway clustering." *Journal of Business & Economic Statistics* 29 (2).
- Costinot, Arnaud, Dave Donaldson, and Cory Smith. 2012. "Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from a 9 Million-Field Partition of the Earth." Tech. rep.
- Cressie, Noel and Christopher K Wikle. 2011. *Statistics for spatio-temporal data*. John Wiley & Sons.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics* 4 (3):66–95.
- Deryugina, T. 2014. "The Role of Transfer Payments in Mitigating Shocks: Evidence from the Impact of Hurricanes." *Working paper* .
- Deschênes, Olivier and Michael Greenstone. 2007. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather." *American Economic Review* 97:354–385.
- . 2011. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics* 3:152–185.
- Deschênes, Olivier, Michael Greenstone, and Jonathan Guryan. 2009. "Climate change and birth weight." *The American Economic Review* :211–217.

- Deschênes, Olivier, Michael Greenstone, and Jonathan Guryan. 2009. "Climate Change and Birth Weight." *American Economic Review: Papers and Proceedings* 99:211–217.
- Easterly, W. and R. Levine. 2003. "Tropics, germs and crops: how endowments influence economic development." *J. of Monetary Econ.* 50:3–39.
- 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." *The American Economic Review* 102 (7):3749–3760.
- Froom, Paul, Yeheskial Caine, Igal Shochat, and Joseph Ribak. 1993. "Heat Stress and Helicopter Pilot Errors." *Journal of Occupational and Environmental Medicine* 35:720–724.
- Gallup, John Luke, Jeffrey D Sachs, and Andrew D Mellinger. 1999. "Geography and economic development." *International regional science review* 22 (2):179–232.
- Graff Zivin, Joshua and Matthew Neidell. 2014. "Temperature and the allocation of time: Implications for climate change." *Journal of Labor Economics* 32 (1):1–26.
- Hartwick, John M. 1977. "Intergenerational equity and the investing of rents from exhaustible resources." *The American Economic Review* :972–974.
- Healy, Andrew and Neil Malhotra. 2009. "Myopic Voters and Natural Disaster Policy." *American Political Science Review* 103 (03):387–406.
- Hornbeck, Richard. 2012. "The Enduring Impact of the American Dust Bowl: Short-and Long-Run Adjustments to Environmental Catastrophe." *The American Economic Review* 102 (4):1477–1507.
- Houser, Trevor, Robert Kopp, Solomon M. Hsiang, Michael Delgado, Amir Jina, Kate Larsen, Michael Mastrandrea, Shashank Mohan, Robert Muir-Wood, DJ Rasmussen, James Rising, and Paul Wilson. forthcoming. *American Climate Prospectus: Economic Risks in the United States*. Columbia University Press.
- Hsiang, Solomon M. 2010. "Temperatures and Cyclones strongly associated with economic production in the Caribbean and Central America." *Proceedings of the National Academy of Sciences* 107 (35):15367–15372.
- Hsiang, Solomon M, Marshall Burke, and Edward Miguel. 2013. "Quantifying the influence of climate on human conflict." *Science* 341 (6151):1235367.
- Hsiang, Solomon M. and Amir S. Jina. 2014. "The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones." Working Paper 20352, National Bureau of Economic Research.
- Hsiang, Solomon M., Kyle C. Meng, and Mark A. Cane. 2011. "Civil conflicts are associated with the global climate." *Nature* 476 (7361):438–441.
- Hsiang, Solomon M. and Daiju Narita. 2012. "Adaptation to Cyclone Risk: Evidence from the Global Cross-Section." *Climate Change Economics* 3 (2).
- Jones, Benjamin and Benjamin Olken. 2010. "Climate shocks and exports." *American Economic Review: Papers and Proceedings* 100:454–459.

- Kahn, Matthew E. 2005. "The Death Toll from Natural Disasters: the role of Income, Geography, and Institutions." *The Review of Economics and Statistics* 87:271–284.
- Key, Nigel, Stacy Sneeringer, and David Marquardt. 2014. "Climate Change, Heat Stress, and US Dairy Production." USDA Economic Research Report Number 175.
- Mackworth, NH. 1947. "High incentives versus hot and humid atmospheres in a physical effort task." *British Journal of Psychology. General Section* 38 (2):90–102.
- Mackworth, Norman H. 1946. "Effects of Heat on Wireless Telegraphy Operators Hearing and Recording Morse Messages." *British Journal of Industrial Medicine* 3:143–158.
- Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw. 1994. "The Impact of Global Warming on Agriculture: A Ricardian Analysis." *The American Economic Review* 84:753–771.
- Nickell, Stephen. 1981. "Biases in dynamic models with fixed effects." *Econometrica* :1417–1426.
- Nordhaus, W. D. 2006. "Geography and macroeconomics: New Data and new findings." *Proceedings of the National Academy of Sciences* 103 (10).
- Nordhaus, William. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.
- Olmstead, Alan L. and Paul W. Rhode. 2011. "Adapting North American wheat production to climatic challenges, 1839–2009." *Proceedings of the National Academy of Sciences* 108 (2):480–485.
- Poterba, James M. 1993. "Global Warming Policy: A Public Finance Perspective." *The Journal of Economic Perspectives* 7 (4):pp. 47–63.
- Roberts, Michael J. and Wolfram Schlenker. 2011. *The Evolution of Heat Tolerance of Corn: Implications for Climate Change*. University of Chicago Press, 225–251.
- Schelling, Thomas C. 1992. "Some Economics of Global Warming." *The American Economic Review* 82 (1):pp. 1–14.
- Schlenker, W., W. M. Hanemann, and A. C. Fisher. 2005. "Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach." *American Economic Review: Papers and Proceedings* 95 (1):395–406.
- Schlenker, Wolfram and Michael 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):15594–15598.
- Schlenker, Wolfram, Michael J Roberts, and David B Lobell. 2013. "US maize adaptability." *Nature Climate Change* 3 (8):690–691.
- Seppanen, Olli, William J Fisk, and QH Lei. 2006. "Effect of temperature on task performance in office environment." Tech. rep., Lawrence Berkeley National Laboratory.
- Solow, Robert. 2012. "Sustainability: an economist's perspective." In *Economics of the Environment*, edited by Robert Stavins. W. W. Norton & Company.

- Stern, Nicholas. 2006. *Stern Review: The Economics of Climate Change*. Cambridge University Press.
- Tol, Richard SJ. 2009. “The economic effects of climate change.” *The Journal of Economic Perspectives* 23 (2):29–51.
- Weitzman, Martin L. 2009. “On modeling and interpreting the economics of catastrophic climate change.” *The Review of Economics and Statistics* 91 (1):1–19.
- Welch, Jarrod R., Jeffrey R. Vincent, Maximilian Auffhammer, Piedad F. Moya, Achim Dobermann, and David Dawe. 2010. “Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures.” *Proceedings of the National Academy of Sciences* .




Figure 1: Example demonstrating the identifying variation in the distribution of county-level daily temperatures from Middlesex county, Massachusetts. The average daily temperature distribution (grey) is absorbed by the county fixed effect. Annual realizations of daily temperature distributions are characterized by the number of days the county falls in each 3°C temperature bin in each year—for example the 2010 distribution for Middlesex is the black-outline overlaid histogram. Idiosyncratic differences between the number of days in each temperature bin for a given year and the climatological average generate the identifying variation used for inference. The daily response function is estimated by computing how changes in the number of days in each bin correspond with total end-of-year income per capita.

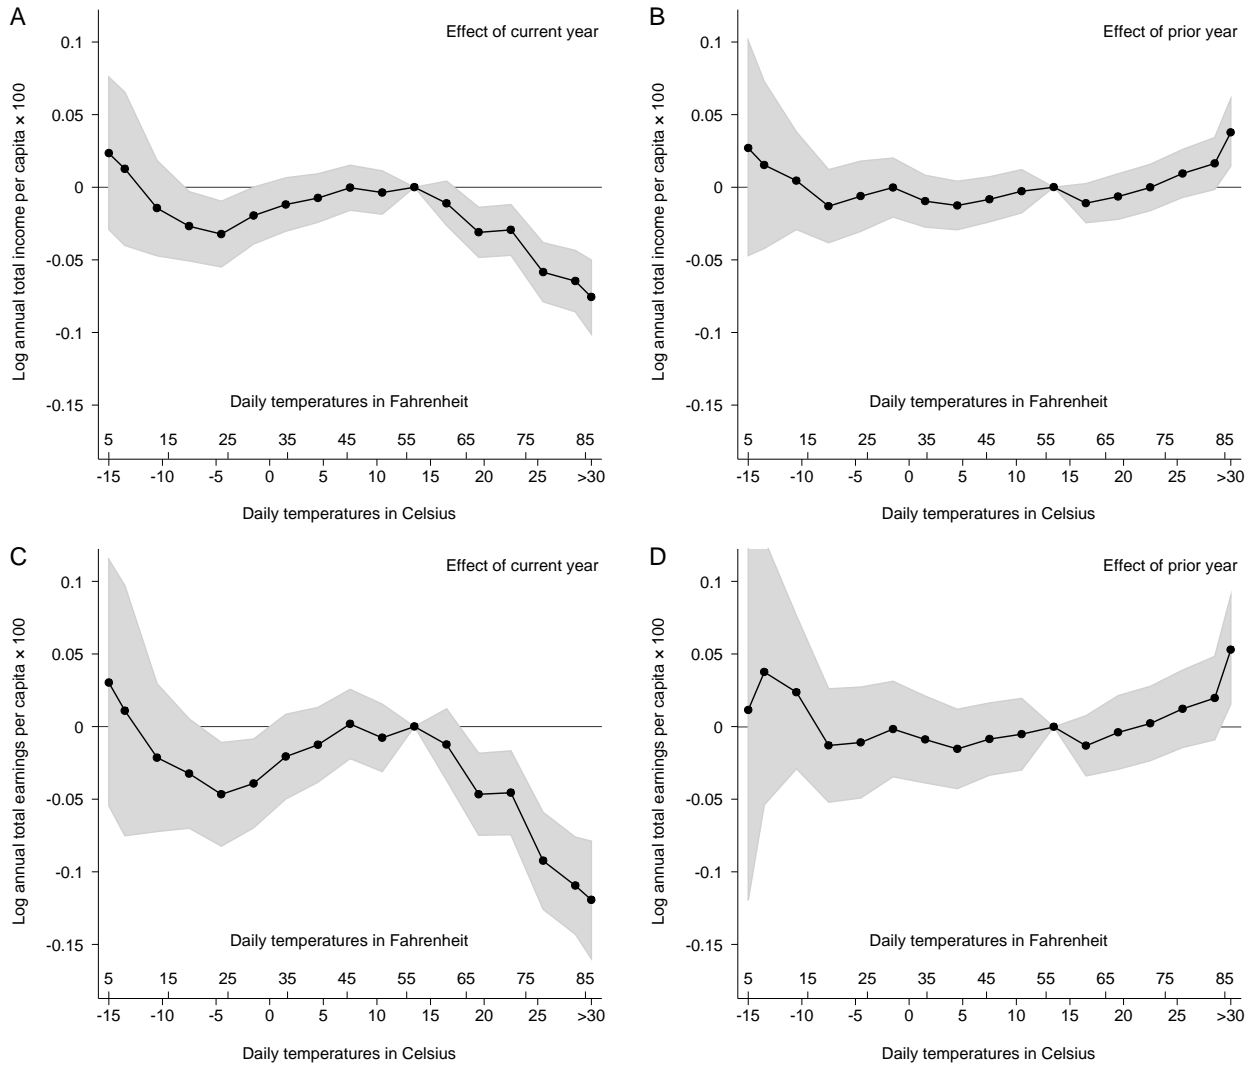


Figure 2: (A) The effect of daily average temperatures on log annual total personal income per capita $\times 100$ (i.e. percentage points) in US counties for 1970-2010. For reference, an average day contributes $\frac{1}{365} = 0.27\%$ of annual income. (B) The effect of daily average temperatures in the prior year on income per capita. (C) Same as Panel A, but for total earnings per capita. (D) Same as Panel B, but for total earnings per capita. Panels A and B are estimated simultaneously in a single regression model. The same is true for Panels C and D.

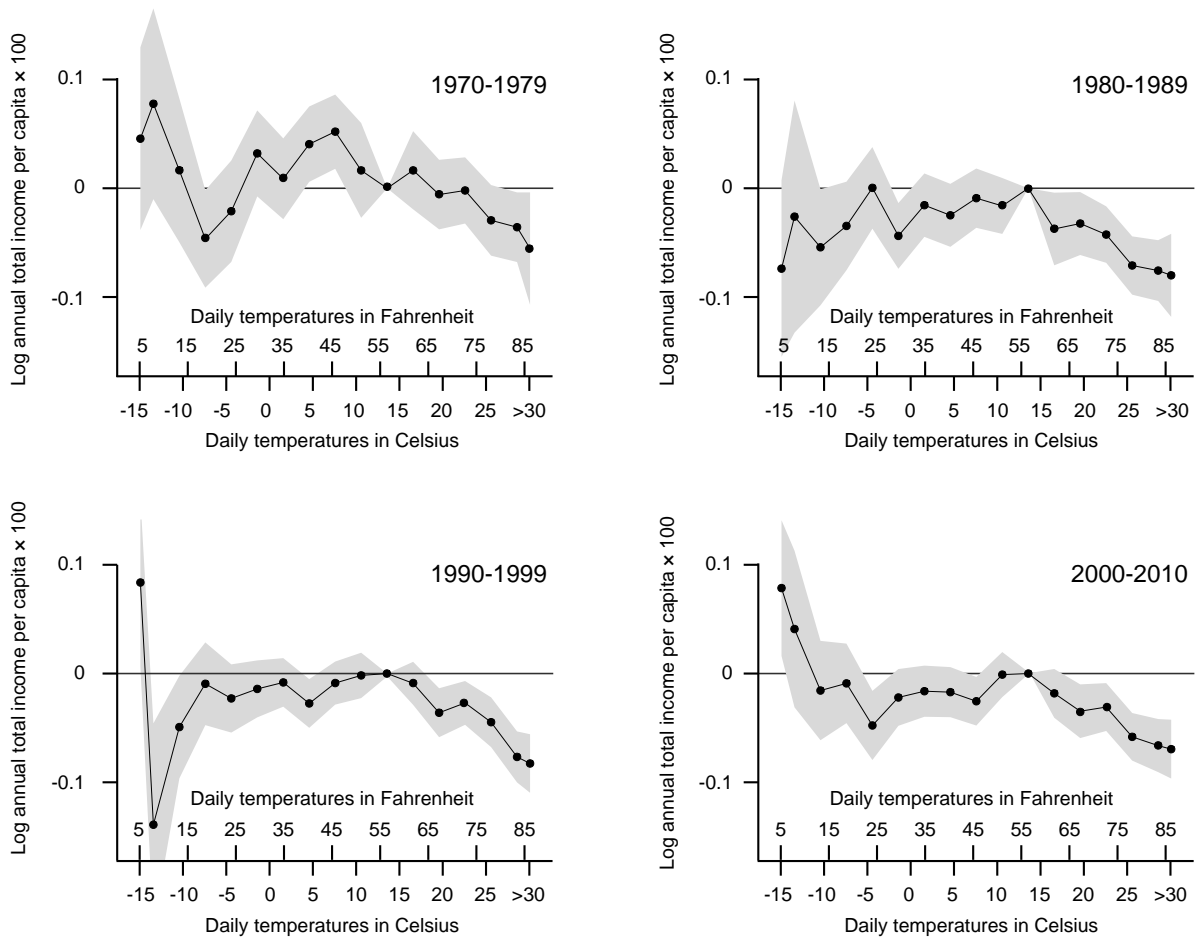


Figure 3: The effect of daily average temperatures on log annual total personal income per capita $\times 100$ (i.e. percentage points) in US counties. Each panel uses a decade-long subsample of the data.




Figure 4: (A) Log personal income per capita in response to daily temperatures on weekdays. (B) The same, but for temperatures on weekend days. Panels A and B are estimated simultaneously in a single regression model.

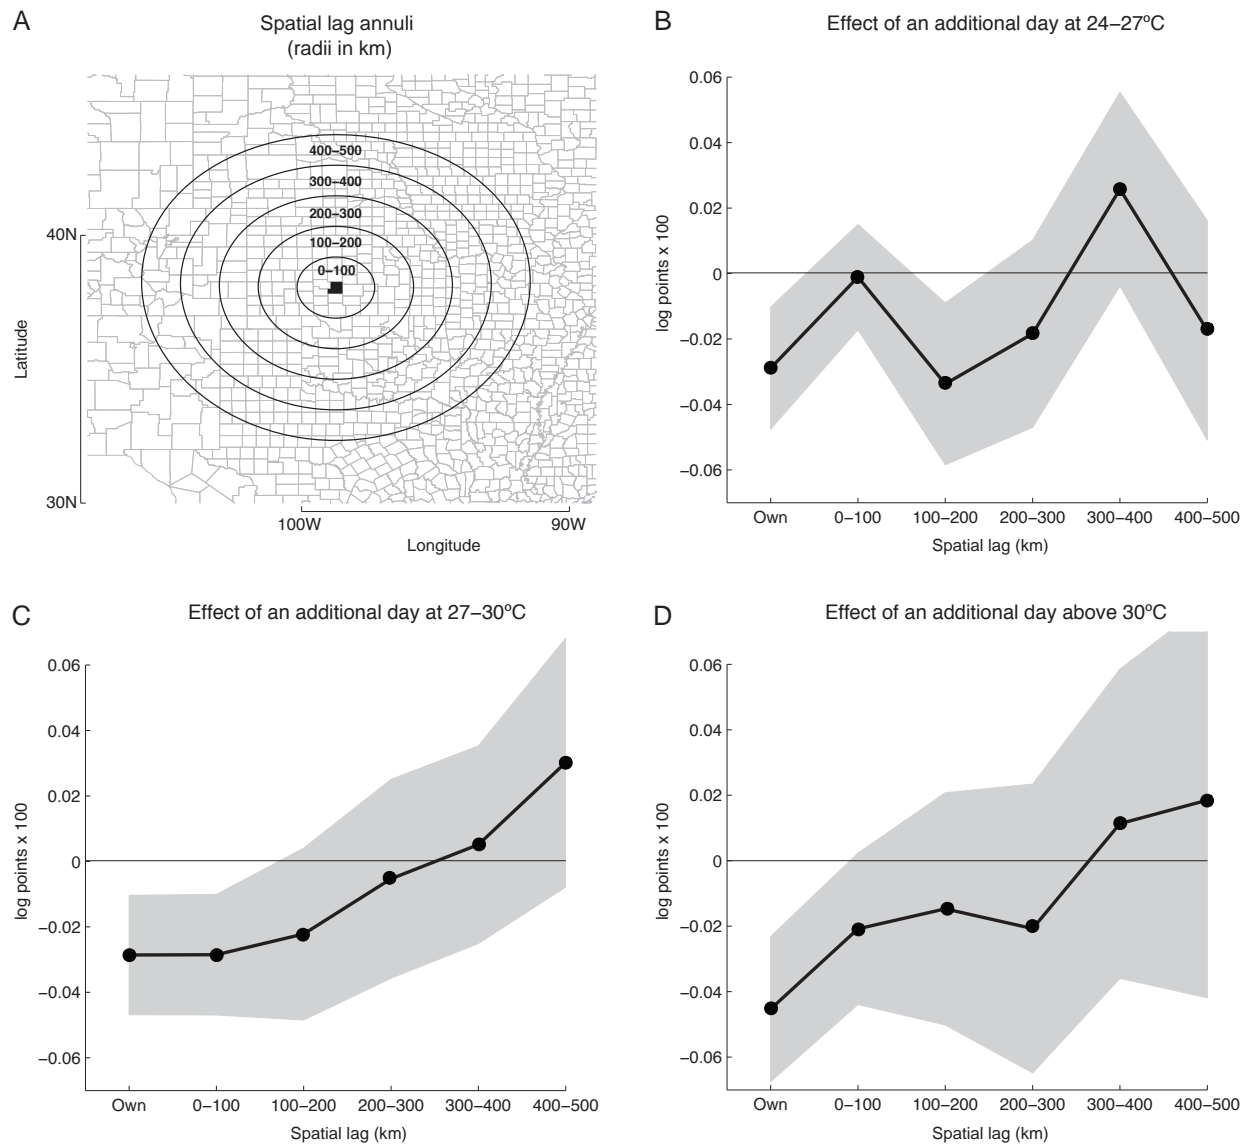


Figure 5: (A) Example of annuli used to construct spatial lags, relative to Stafford, Kansas (black). (B) Effect on i of each additional day at 24–27°C for i and 24–27°C days experienced by j 's at various distances from i . (C) Same but for 27–30°C. (D) Same but for > 30°C. All effects in (B)–(D) are estimated simultaneously, along with own effects for lower temperatures and all controls.

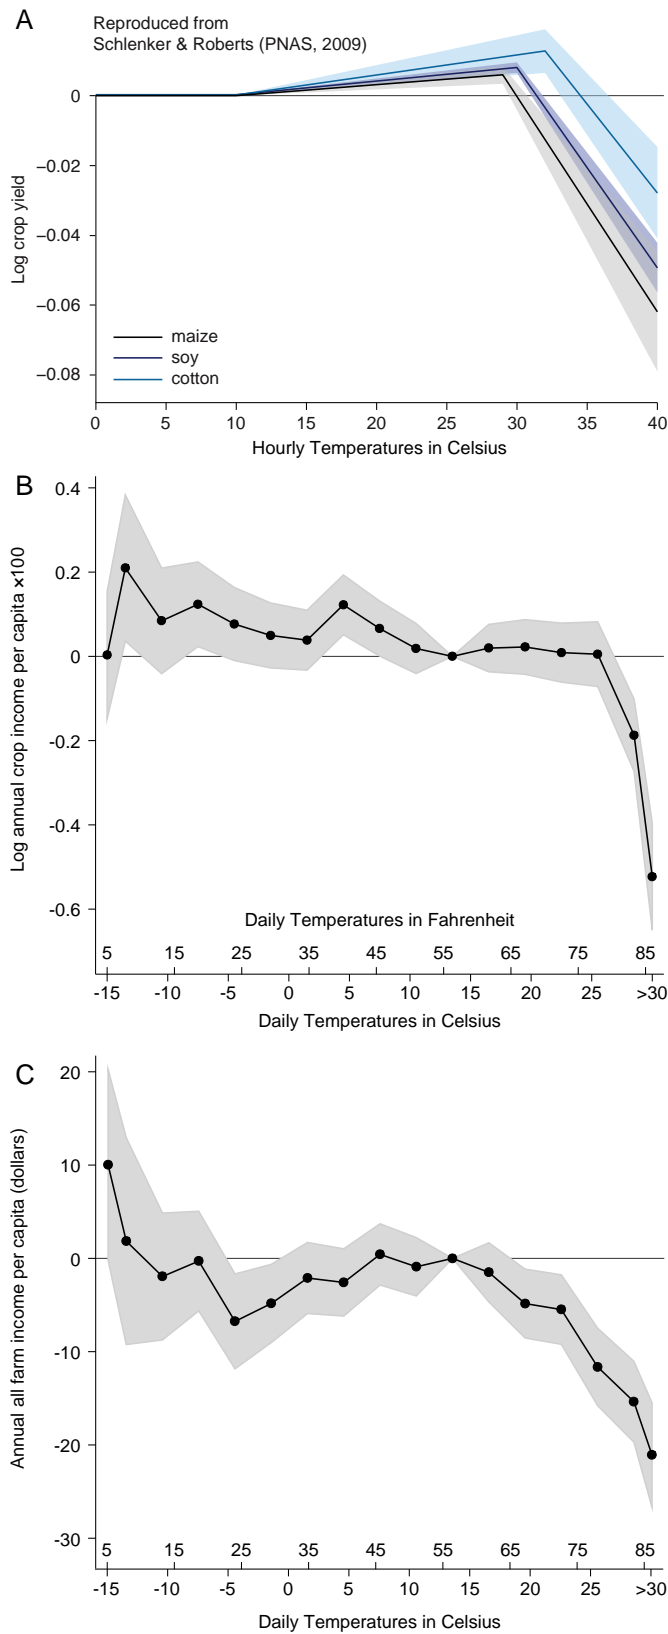


Figure 6: (A) Log annual crop yields vs. temperature during growing season for three major crops, reproduced from Schlenker and Roberts (2009). Yield effects are depicted as the effect of 24 hours at exact temperatures. (B) The effect of daily average temperature on log income from crops per capita, from this study. (C) The effect of daily average temperature on farm income per capita (in levels), from this study.

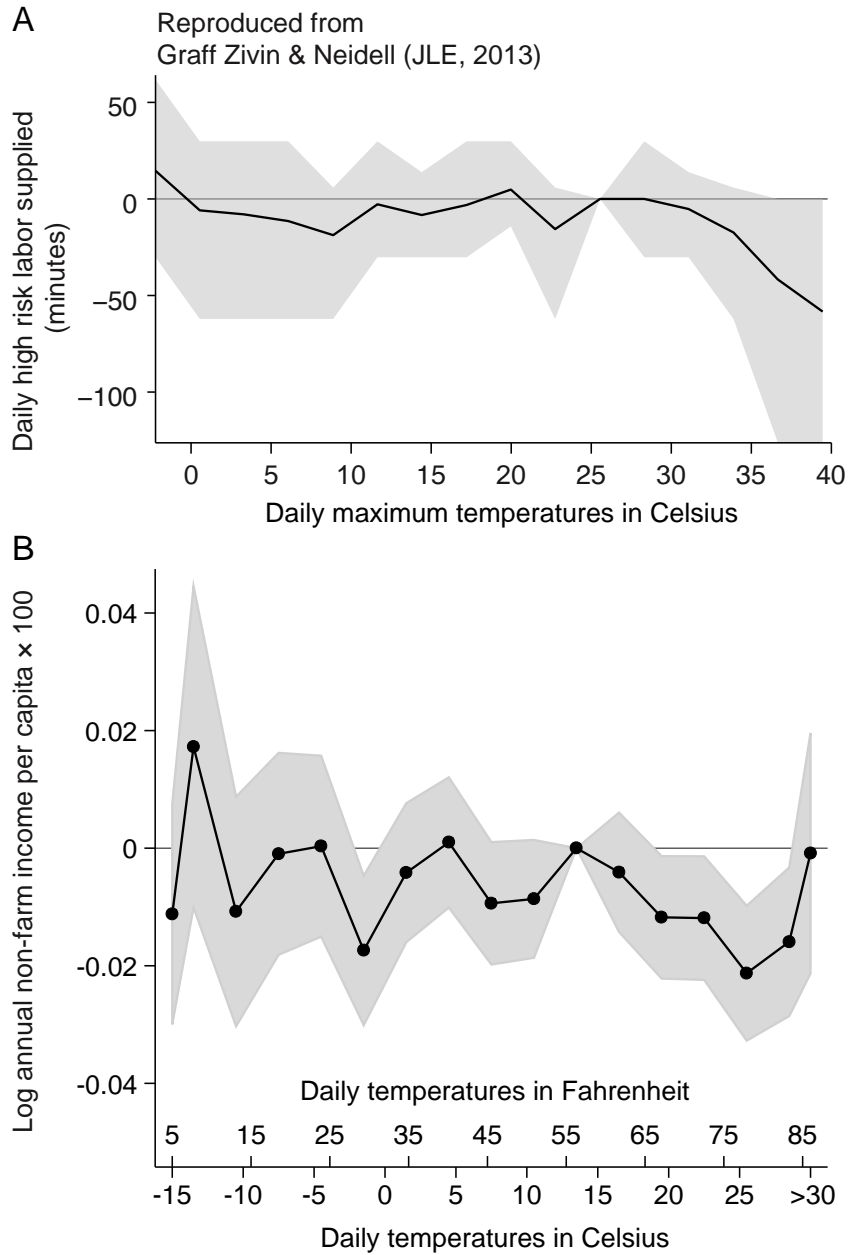


Figure 7: (A) Change in minutes of labor supplied per day for high-risk workers vs daily maximum temperature, reproduced from Graff Zivin and Neidell (2014). High-risk workers are defined as workers who are likely exposed to outdoor temperatures (includes manufacturing). (B) The effect of daily average temperature on non-farm income per capita, from this study.

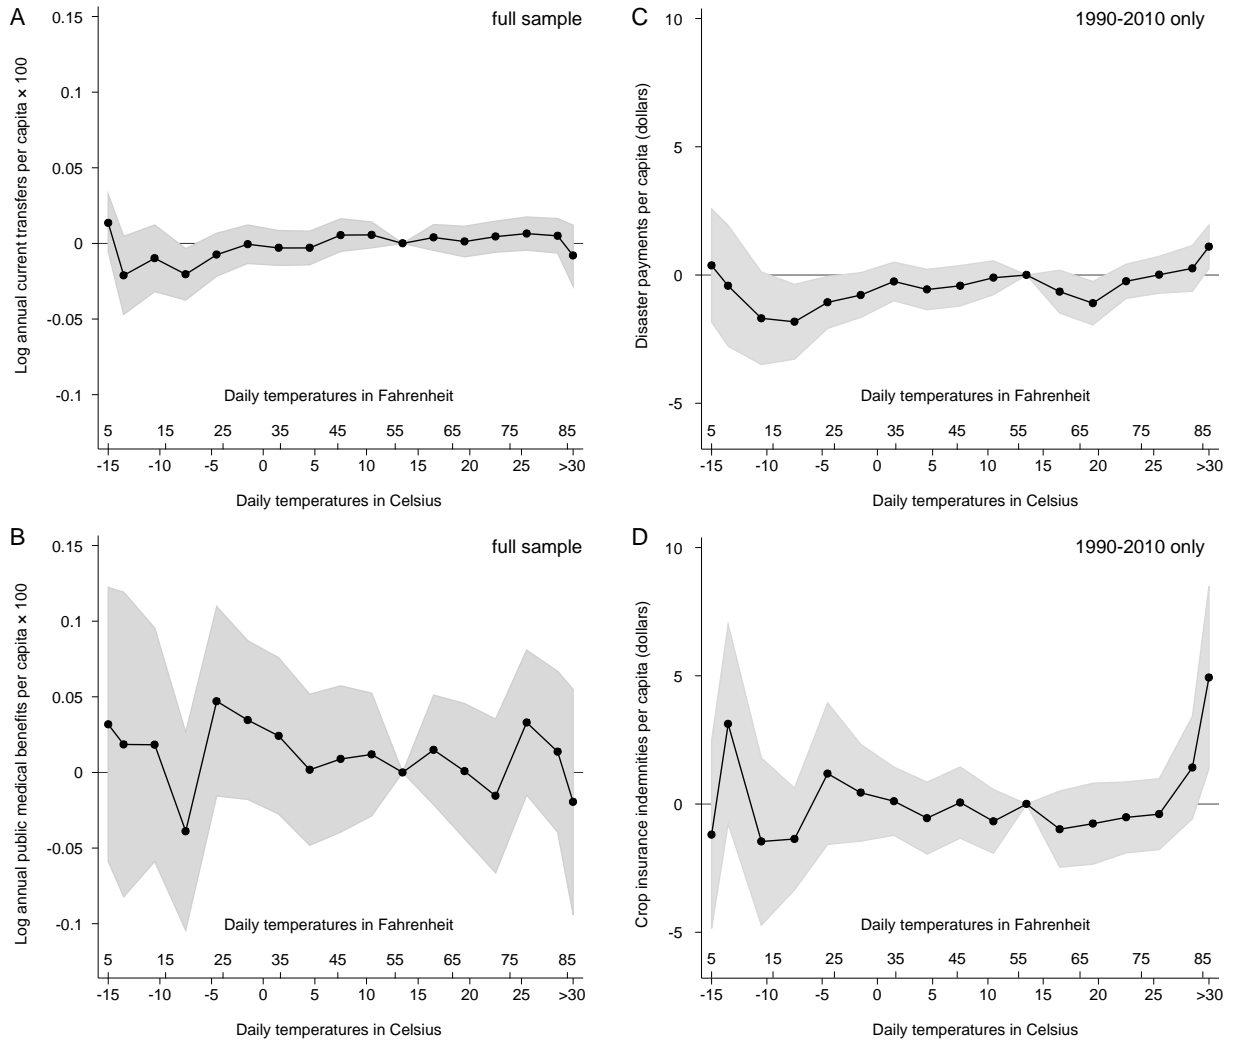


Figure 8: (A) The effect of daily average temperatures on log total government transfers to individuals per capita. (B) The same as A, but for the subset of transfers that are public medical benefits. (C) The effect of daily temperatures on *ad hoc* crop disaster payments per capita (in levels). (D) The effect of daily temperatures on crop insurance indemnities per capita.




Figure 9: (A) Comparison of observed nationally aggregated total personal income (dashed black) with model prediction using observed temperature (orange) and temperature that is optimized to maximize income (blue), i.e. 12-15°C every county-day. Shaded regions are 95% confidence intervals from bootstrapped estimates. (B) Distribution of bootstrapped income growth rates implied by Panel A.

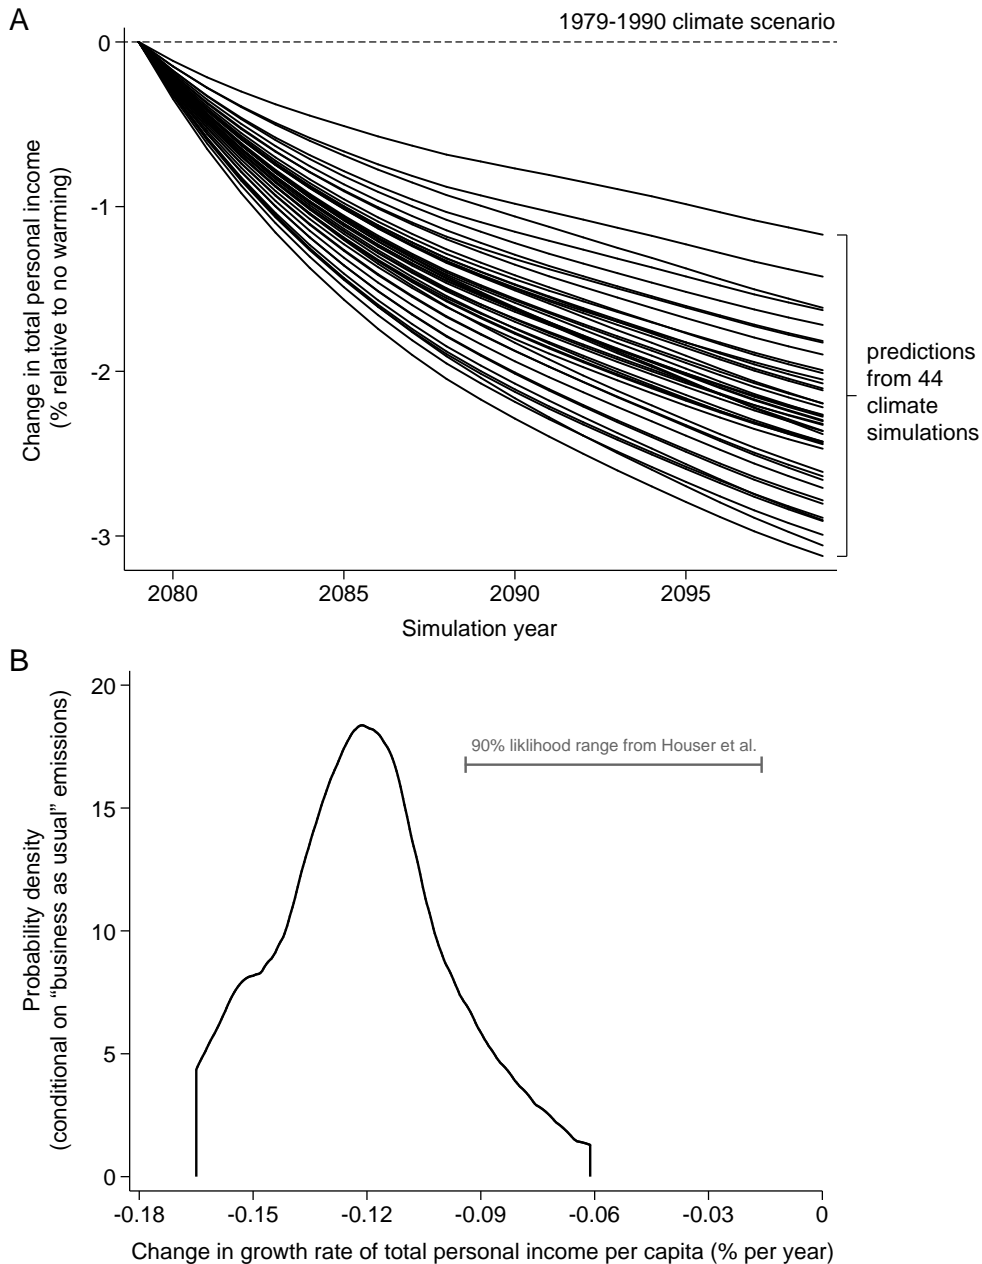


Figure 10: (A) Estimated change in total personal income in a “business as usual” climate change scenario (RCP 8.5) relative to a projection where the climate is fixed to 1970-1990 conditions. Both simulations are initialized in 2080 with the same income level. Each line is one of 44 different climate model projections using the same emissions scenario, from Houser et al. (forthcoming). (B) Estimated distribution for the growth rate effect of “business as usual” climate change scenarios shown in Panel A. Following Houser et al., each of the 44 model projections is weighted to reflect the overall probability distribution of global climate sensitivities estimated from historical climate data. 90% likelihood range for analogous growth rate change in Houser et al. shown in gray for comparison. Estimates from Houser et al. are computed by integrating detailed county-level damage projections for labor supply, agriculture, mortality, energy, cyclones, and sea level rise into a computable general equilibrium model.

Table 1: Summary statistics

	(1) Mean	(2) Std. Dev.	(3) Min	(4) Max	(5) Obs
Population	111,722	341,233	209	9,889,056	76,646
Personal income per capita	26,806	8,438	6,356	136,936	76,646
Non-farm personal income per capita	16,710	10,549	2,738	336,356	76,646
Percent of personal income that is non-farm income	61.67	27.53	8	916	76,646
Percent of personal income that is wage/salary income	45.16	22.59	9	757	76,646
Percent of personal income that is farm income	5.24	8.66	-235	77	76,646
Percent of personal income that is rents	18.15	5.93	2	123	76,646
Percent of personal income that is transfers	16.88	6.38	2	65	76,646

Source: Regional Economic Information Systems. Unit of observation is a county-year. All monetary amounts are in 2011 dollars.

Table 2: The effect of daily temperature on log total annual income per capita

		(1)	(2)	(3)	(4a)	(4b)
Daily temperature		Pooled week model			Joint weekend-weekday model	
(Celsius)	(Fahrenheit)	All days	All days	All days	Weekday	Weekend
>30°C	>86°F	-0.000757*** (0.00013)	-0.000658*** (0.000235)	-0.000279** (0.000129)	-0.001255*** (0.000214)	0.000496*** (0.000192)
27-30°C	80.6-86°F	-0.000646*** (0.000108)	-0.000436** (0.00017)	-0.000235** (0.000097)	-0.000874*** (0.000144)	0.000262** (0.000116)
24-27°C	75.2-80.6°F	-0.000585*** (0.000104)	-0.000866*** (0.000159)	-0.000109 (0.000091)	-0.000666*** (0.00013)	0.000092 (0.000098)
21-24°C	69.8-75.2°F	-0.000294*** (0.000089)	-0.000445*** (0.000137)	0.000033 (0.00008)	-0.000407*** (0.000109)	0.00012 (0.000079)
18-21°C	64.4-69.8°F	-0.000311*** (0.000088)	-0.000672*** (0.000138)	-0.000012 (0.000075)	-0.000390*** (0.000099)	0.000094 (0.000067)
15-18°C	59-64.4°F	-0.000111 (0.000078)	-0.000301** (0.000121)	0.000017 (0.000062)	-0.000169** (0.000076)	0.000076 (0.000047)
12-15°C	53.6-59°F	0 -	0 -	0 -	0 -	0 -
9-12°C	48.2-53.6°F	-0.000036 (0.000076)	-0.000288** (0.000124)	0.000004 (0.000059)	0.000017 (0.000064)	-0.000054 (0.000049)
6-9°C	42.8-48.2°F	-0.000003 (0.000079)	-0.00017 (0.000135)	-0.000001 (0.000069)	-0.000054 (0.000078)	0.000054 (0.000065)
3-6°C	37.4-42.8°F	-0.000074 (0.000085)	-0.000195 (0.000149)	0.00004 (0.00008)	-0.000071 (0.000085)	0.000014 (0.000077)
0-3°C	32-37.4°F	-0.000119 (0.000094)	-0.000271* (0.000157)	0.000123 (0.000084)	-0.000038 (0.000098)	-0.000063 (0.000086)
-3-0°C	26.6-32°F	-0.000195* (0.0001)	-0.000420** (0.00017)	-0.000024 (0.00009)	-0.000064 (0.000113)	-0.000114 (0.000096)
-6--3°C	21.2-26.6°F	-0.000322*** (0.000115)	-0.000657*** (0.000187)	-0.000195* (0.000105)	-0.000081 (0.000139)	-0.000229** (0.000114)
-9--6°C	15.8-21.2°F	-0.000269** (0.000122)	-0.000616*** (0.000216)	0.000007 (0.000115)	0.000032 (0.000157)	-0.000288** (0.000138)
-12--9°C	10.4-15.8°F	-0.000144 (0.000167)	-0.000149 (0.00025)	0.000270** (0.000135)	0.000237 (0.000182)	-0.000364** (0.000153)
-15--12°C	5-10.4°F	0.000126 (0.000269)	0.000691* (0.000376)	0.000083 (0.000179)	0.000666** (0.000268)	-0.000534** (0.000227)
< -15°C	<5°F	0.000234 (0.000268)	-0.000162 (0.000331)	0.000478*** (0.000182)	0.000373 (0.000281)	-0.000125 (0.000278)
Observations		76,576	78,635	76,576	76,576	76,576
R-squared		0.93	0.85	0.95	0.93	0.93
County Fixed Effects		Y	Y	Y	Y	Y
Lagged Dependant Variable		Y	No	Y	Y	Y
Year Fixed Effects		Y	Y	No	Y	Y
State-by-Year Fixed Effects		No	No	Y	No	No


Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Standard errors (in parentheses) clustered by county and by state-year. All outcomes are in dollars per capita. Controls include year and county fixed effects, lagged weather variables and the lagged dependent variable. Omitted category is 12-15 degrees Celsius (53.6-59 degrees Fahrenheit). Columns 4a and 4b are estimates simultaneously as part of the same model.

Table 3: Estimating the effect of daily temperature on annual county income per capita in dollars

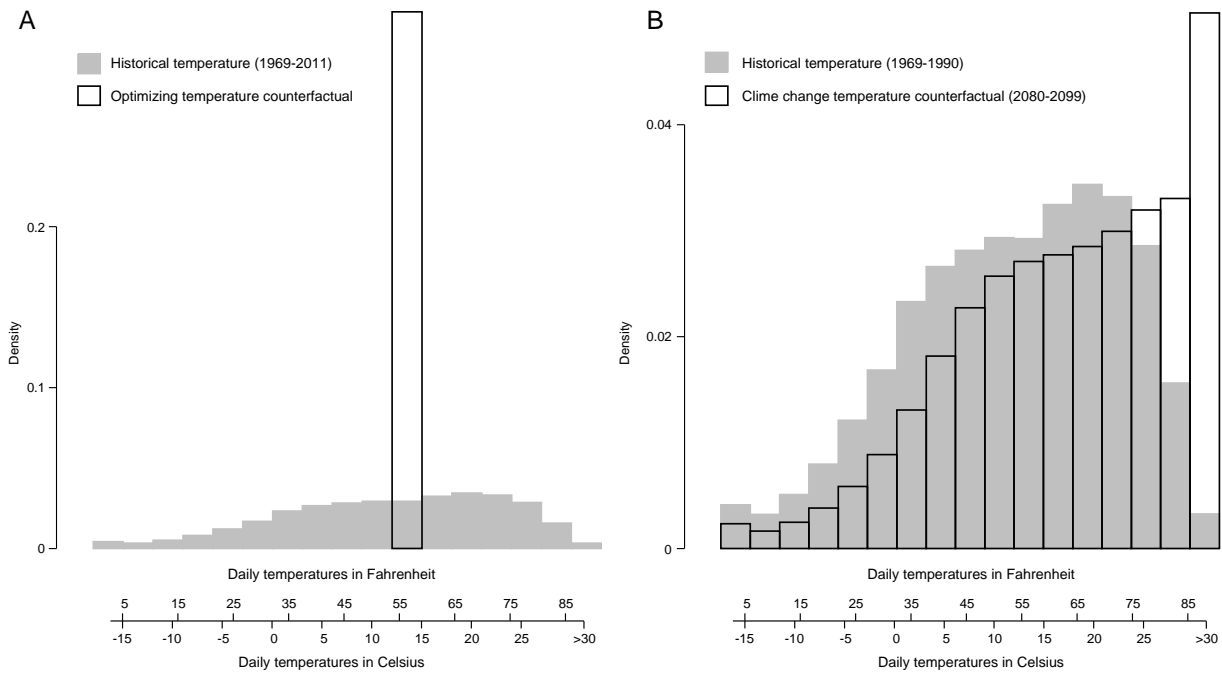
		(1)	(2)	(3)	(4)	(5)	(6)
Daily temperature		Unweighted income			Population-weighted income		
(Celsius)	(Fahrenheit)	Total	Farm	Non-farm	Total	Farm	Non-farm
>30°C	>86°F	-20.56*** (3.33)	-21.07*** (2.85)	0.53 (1.81)	-2.39 (1.74)	-2.87*** (0.60)	0.30 (2.08)
27-30°C	80.6-86°F	-16.71*** (2.73)	-15.36*** (2.21)	-2.27** (1.06)	-1.48 (2.35)	-2.08*** (0.41)	-1.95 (1.85)
24-27°C	75.2-80.6°F	-14.78*** (2.61)	-11.63*** (2.12)	-3.03*** (0.94)	-4.80*** (1.66)	-0.67** (0.34)	-4.05*** (1.52)
21-24°C	69.8-75.2°F	-6.92*** (2.29)	-5.49*** (1.90)	-1.80** (0.89)	-2.72* (1.59)	-0.11 (0.34)	-3.95** (1.55)
18-21°C	64.4-69.8°F	-7.46*** (2.30)	-4.85** (1.88)	-1.75** (0.87)	-2.12 (1.56)	0.09 (0.31)	-2.64* (1.59)
15-18°C	59-64.4°F	-2.77 (1.99)	-1.50 (1.61)	-0.50 (0.87)	-1.06 (1.70)	0.24 (0.35)	-2.35* (1.42)
12-15°C	53.6-59°F	0 -	0 -	0 -	0 -	0 -	0 -
9-12°C	48.2-53.6°F	-1.00 (2.01)	-0.90 (1.59)	-0.73 (0.85)	-0.21 (1.85)	-0.55 (0.44)	-1.24 (1.67)
6-9°C	42.8-48.2°F	-1.07 (2.11)	0.42 (1.67)	-1.44 (0.91)	-1.91 (1.35)	-0.37 (0.30)	-1.97 (1.36)
3-6°C	37.4-42.8°F	-2.47 (2.37)	-2.58 (1.84)	0.18 (0.98)	-1.66 (1.69)	-0.40 (0.44)	-2.82* (1.68)
0-3°C	32-37.4°F	-3.12 (2.47)	-2.11 (1.94)	-0.16 (1.02)	1.07 (1.72)	-0.49 (0.38)	0.26 (1.69)
-3-0°C	26.6-32°F	-5.47** (2.66)	-4.84** (2.14)	-2.89*** (1.12)	-0.48 (1.90)	-0.69 (0.44)	-2.52 (1.91)
-6--3°C	21.2-26.6°F	-10.61*** (3.13)	-6.76*** (2.60)	-0.61 (1.38)	-4.09* (2.26)	-0.78 (0.52)	-5.34*** (2.31)
-9--6°C	15.8-21.2°F	-5.74* (3.32)	-0.29 (2.73)	0.49 (1.57)	-0.37 (2.50)	-0.26 (0.53)	0.88 (2.34)
-12--9°C	10.4-15.8°F	-1.73 (4.38)	-1.94 (3.47)	-0.99 (1.77)	-3.70 (3.94)	-1.20* (0.66)	-3.09 (3.81)
-15--12°C	5-10.4°F	4.28 (6.74)	1.85 (5.66)	2.32 (2.40)	0.54 (4.90)	-0.85 (1.36)	2.24 (4.92)
< -15°C	<5°F	6.55 (6.85)	10.03* (5.22)	-2.02 (1.57)	1.57 (3.71)	4.05** (1.84)	-3.66 (2.63)
Observations		76,576	75,212	76,576	76,576	75,212	76,576
R-squared		0.93	0.34	0.95	0.99	0.32	0.98

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Standard errors (in parentheses) clustered by county and by state-year. All outcomes are in 2011 US dollars per capita and are Winsorized at the 1st and 99th percentile. Controls include year and county fixed effects, lagged weather variables and the lagged dependent variable. Omitted category is 12-15 degrees Celsius (53.6-59 degrees Fahrenheit).

Appendix



Appendix Figure A1: Number of years that each county has a complete record of daily average temperatures and daily rainfall. Years with incomplete records are dropped from the sample.



Appendix Figure A2: (A) Distribution of daily temperatures in the sample and used in the historical simulation (grey) and distribution of daily temperatures used in the counterfactual simulation where counties select an optimal temperature that maximizes daily production (black outline). (B) Distribution of daily temperatures used in the historical simulation (grey) and distribution of daily temperatures used in the “business as usual” climate change simulation (black outline). Histogram is weighted distribution across 44 climate models, data from Houser et al. (forthcoming).

Appendix Table A1: The effect of daily temperature on components of county income

		(1)	(2)	(3)	(4)	(5)
Daily temperature		log	log	log	log	log
(Celsius)	(Fahrenheit)	total income per capita	net earnings per capita	non-farm inc. per capita	crop income per capita	transfers per capita
>30°C	>86°F	-0.000757*** (0.00013)	-0.001193*** (0.000207)	-0.000009 (0.000104)	-0.005230*** (0.000647)	-0.000081 (0.000103)
27-30°C	80.6-86°F	-0.000646*** (0.000108)	-0.001095*** (0.000171)	-0.000159** (0.000064)	-0.001874*** (0.000438)	0.00005 (0.000058)
24-27°C	75.2-80.6°F	-0.000585*** (0.000104)	-0.000924*** (0.00017)	-0.000213*** (0.000059)	0.000051 (0.000392)	0.000065 (0.000056)
21-24°C	69.8-75.2°F	-0.000294*** (0.000089)	-0.000456*** (0.000147)	-0.000119** (0.000054)	0.000088 (0.000358)	0.000045 (0.000052)
18-21°C	64.4-69.8°F	-0.000311*** (0.000088)	-0.000466*** (0.000144)	-0.000118** (0.000053)	0.000223 (0.000331)	0.000012 (0.000052)
15-18°C	59-64.4°F	-0.000111 (0.000078)	-0.000124 (0.000126)	-0.000041 (0.000052)	0.000196 (0.000287)	0.000039 (0.000044)
12-15°C	53.6-59°F	0 -	0 -	0 -	0 -	0 -
9-12°C	48.2-53.6°F	-0.000036 (0.000076)	-0.000078 (0.000118)	-0.000086* (0.000051)	0.000187 (0.000303)	0.000056 (0.000044)
6-9°C	42.8-48.2°F	-0.000003 (0.000079)	0.000018 (0.000122)	-0.000094* (0.000053)	0.000659** (0.000333)	0.000055 (0.000055)
3-6°C	37.4-42.8°F	-0.000074 (0.000085)	-0.000126 (0.000131)	0.00001 (0.000056)	0.001223*** (0.000362)	-0.00003 (0.000057)
0-3°C	32-37.4°F	-0.000119 (0.000094)	-0.000207 (0.000149)	-0.000042 (0.00006)	0.000384 (0.000363)	-0.00003 (0.000059)
-3-0°C	26.6-32°F	-0.000195* (0.0001)	-0.000392** (0.000156)	-0.000174*** (0.000065)	0.000497 (0.000393)	-0.000005 (0.000065)
-6--3°C	21.2-26.6°F	-0.000322*** (0.000115)	-0.000467** (0.000182)	0.000003 (0.000079)	0.000768* (0.000443)	-0.000075 (0.000072)
-9--6°C	15.8-21.2°F	-0.000269** (0.000122)	-0.000324* (0.000192)	-0.000001 (0.000088)	0.001236** (0.000513)	-0.000204** (0.000087)
-12--9°C	10.4-15.8°F	-0.000144 (0.000167)	-0.000214 (0.00026)	-0.000108 (0.0001)	0.000843 (0.00064)	-0.000098 (0.000112)
-15--12°C	5-10.4°F	0.000126 (0.000269)	0.000109 (0.000439)	0.000172 (0.000139)	0.002097** (0.000885)	-0.000211 (0.000132)
< -15°C	<5°F	0.000234 (0.000268)	0.000303 (0.000433)	-0.000112 (0.000096)	0.00003 (0.00077)	0.000136 (0.000097)
Observations		76,576	76,571	76,576	75,312	76,574
R-squared		0.93	0.80	0.95	0.73	0.99

Significance levels: * 10 percent, ** 5 percent, *** 1 percent. Standard errors (in parentheses) clustered by county and by state-year. All outcomes are in dollars per capita. Controls include year and county fixed effects, lagged weather variables and the lagged dependent variable. Omitted category is 12-15 degrees Celsius (53.6-59 degrees Fahrenheit).