

NBER WORKING PAPER SERIES

THE MARGINAL PRODUCT OF CLIMATE

Tatyana Deryugina
Solomon Hsiang

Working Paper 24072
<http://www.nber.org/papers/w24072>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2017

Portions of this analysis were previously circulated as an NBER working paper under the title "Does the Environment Still Matter? Daily Temperature and Income in the United States." We thank Max Auffhammer, Marshall Burke, Tamma Carleton, Melissa Dell, Olivier Deschenes, Don Fullerton, Michael Greenstone, Hilary Hoynes, Amir Jina, Larry Karp, Bentley Macleod, Kyle Meng, Edward Miguel, Robert Pindyck, Billy Pizer, Jonathan Proctor, Julian Reif, James Rising, Michael Roberts, Francois Salanie, Wolfram Schlenker, Richard Schmalensee, Joe Shapiro, Tony Smith, David Strömberg, James Stock, Anna Tompsett, Reed Walker, Frank Wolak, and seminar participants at Chicago, FEEM, IZA, LSE, MIT, NBER EEE Summer Institute, Princeton, Stanford, Stockholm, Toulouse School of Economics, UC Berkeley, UC Davis, UC Santa Barbara, UI Urbana-Champaign, UNC, and Yale for discussions and comments. We thank Tamma Carleton, Michael Delgado, DJ Rasmussen, and Andrew Wilson 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Tatyana Deryugina and Solomon 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.

The Marginal Product of Climate
Tatyana Deryugina and Solomon Hsiang
NBER Working Paper No. 24072
November 2017
JEL No. E23,O4,Q5,R13

ABSTRACT

We develop an empirical approach to value changes to a climate in terms of total market output given optimal factor allocations in general equilibrium. Our approach accounts for unobservable heterogeneity across locations as well as the costs and benefits of adaptation in climates of arbitrary dimension. Importantly, we demonstrate that the Envelope Theorem implies the marginal product of a long-run climate can be exactly identified using only idiosyncratic weather variation. We apply this method to the temperature climate of the modern United States and find that, despite evidence that populations adapt, the marginal product of temperature has remained unchanged during 1970-2010, with high temperatures having low net value. Integrating marginal products recovers a value function for temperature, describing the causal effect of non-marginal climate changes net of adaptive re-optimization. We use this value function to consider the influence of temperature in the current cross-section and a future climate change scenario.

Tatyana Deryugina
Department of Finance
University of Illinois at Urbana-Champaign
515 East Gregory Drive, MC-520
Champaign, IL 61820
and NBER
deryugin@illinois.edu

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

1 Introduction

We consider the contribution of climate to the total market output of United States (US) counties. Empirical construction of these values is challenging because it requires simultaneously accounting for all unobservable differences between counties (Deschênes and Greenstone, 2007) and all endogenous adaptations to climate net of costs (Schlenker, Roberts, and Lobell, 2013)—criteria that no prior approach has yet delivered (Hsiang, 2016). Here we develop a single unified framework that satisfies both of these criteria for an arbitrary number of unobserved allocative adaptations in an economy at general equilibrium responding to a high-dimensional climate. We show that our approach can recover the marginal product of climate by using only idiosyncratic weather variation and a reduced-form estimator.

It is well known that the short-run income response to weather and the long-run income response to climate are not the same mathematical object, due to possible adaptive re-optimization (Mendelsohn, Nordhaus, and Shaw, 1994; Kelly, Kolstad, and Mitchell, 2005). Our core insight is that the *marginal* effect of weather, suitably defined, is nonetheless exactly equal to the *marginal* effect of climate on total output when an economy is in competitive equilibrium. This equality holds locally because the two surfaces that describe the weather-income and climate-income relationships are exactly tangent in the neighborhood of the equilibrium allocation. The surfaces are tangent because no adaptive reallocation occurs due to short-run idiosyncratic weather fluctuations, by definition, and because the effect of any marginal adaptation on income in the long-run is exactly zero as a result of the Envelope Theorem (e.g. Guo and Costello, 2013). Thus, in neither the weather-income nor the climate-income case are local marginal effects influenced by adaptive behavior.

The equivalence between the marginal effects of climate and the marginal effects of weather enables us to measure the local gradient of the long-run relationship implicitly by directly measuring the local gradient of the short-run relationship. The short-run response can be empirically identified by exploiting idiosyncratic variations in weather over time within each location, enabling the removal of unobserved cross-sectional heterogeneity from parameter estimates (Deschênes and Greenstone, 2007). A large number of local gradient estimates observed at “nearby” baseline climates can then be “pieced together” through integration to reconstruct the long-run income response surface, which cannot otherwise be observed directly. The key data requirements necessary for this approach to be valid is a large panel of similar economies that (i) span the *space of possible climates*, which we define precisely below, (ii) are sufficiently densely packed in this space such that integration between positions is reasonable, and (iii) experience short-run weather disturbances that are not “too large” in the sense that they do not perturb the economy so far from its equilibrium that the Envelope Theorem no longer applies.

We demonstrate the application of this approach using a large number of ‘small macro-

economies’ represented by the panel of modern US counties, which plausibly satisfy the criteria above. We examine counties’ local income response to small perturbations in the annual distribution of daily temperatures and discover a remarkably strong and stable relationship between temperatures and production across space, seasons, and time. Importantly, we show that investments in human-made capital—in the form of air conditioning and cities—appear to be partial substitutes for climate in production, in the sense of Hartwick (1977) and Solow (1991). However, we continue to observe large contributions of climate to output even in extremely urbanized contexts and in the twenty-first century, indicating a high net value of certain climates despite the existence of numerous possible margins of adaptation.

Our approach enables us to integrate marginal effects of climate to compute causal effects on production due to non-marginal climate changes. Thus, we can compute differences in economic production that are attributable to differences in contemporaneous locations’ climates, as well as to estimate changes in future production due to projected future warming, holding technology constant. In both cases, our estimates account for both the costs and benefits of all margins of endogenous adaptation. On net, we find that existing climate differences between counties generate substantial differences in their output, with hotter climates having lower average production, *ceteris paribus*, a result that we recover without exploiting cross-sectional variation in estimation of our parameters. For example, we estimate that the climate of Northern Minnesota returns over \$2,000 per capita more annually than the climate of Southern Texas. For similar reasons, projections of future output are substantially reduced once future warming and resulting adaptation are both accounted for. Net of all currently available adaptation technologies, we estimate the value of projected changes in US production due to warming during the twenty-first century at -\$6.7 trillion in net present value (US 2011 dollars, RCP8.5, discounted at 3% annually) in the median scenario, using our preferred specification. The 90% confidence interval of this loss is \$4.7-10.4 trillion when accounting for climate model uncertainty (Burke et al., 2015).¹

An ancillary but potentially important insight from our empirical work is that accounting for adaptive reallocations *increases* total projected losses under warming relative to an approach that assumes uniform marginal effects everywhere. This result is counter to the widely referenced “folk theorem” that accounting for endogenous adaptation empirically should necessarily reduce the estimated impacts of climate change. In our context, we find the marginal damages from warming—which are larger for cooler and less adapted northern counties—are positively correlated across space with the distribution of economic activity. Thus, models that assume uniform marginal effects *under-estimate* future losses because

¹It should be noted that these values do not represent welfare calculations, nor do they account for non-market impacts of warming. They also do not include effects of climate change through non-temperature channels (e.g., floods), and they clearly cannot account for possible future technological innovations that do not yet exist in our data.

marginal damage estimates for northern locations are biased toward zero by pooling them, in estimation, with hotter and more adapted (but less productive) southern counties.

The structure of the paper is as follows. In Section 2, we introduce definitions for climate, the space of all possible climates, the marginal product of climate, the role of climate in a market equilibrium, and the relationship between climate and weather. In Section 3 we derive how the marginal product of climate can be estimated using weather variation and how these estimates can be used to compute non-marginal effects of climate. In Section 4 we explain our empirical implementation in the modern US, deriving how we recover the marginal product of high-dimensional daily temperature distributions. In Section 5 we examine the structure of the marginal product of temperature in the US, including its stability over time, the dynamics of income growth, and the effects of adaptation. In Section 6 we examine what mechanisms might be responsible for these results, considering both different sectors of production (e.g. agriculture, manufacturing) and the role of human-made capital (e.g. air-conditioning) as partial substitutes for climate in production. In Section 7 we use our results to compute the non-marginal effects of temperature on the current cross-section of income and the projected value of future warming. Section 8 discusses important caveats of our analysis and points towards areas for future research.

2 Theoretical setup and statement of the problem

To quantify the full marginal product of a climate, we must compute the economic value generated by an economy facing that climate relative to the same economy when it faces a slightly different climate, including the costs and benefits of any adaptation measures taken in response. Accounting for the benefits of adaptation is relatively straightforward, since the reduced sensitivity of populations is directly observable (e.g. Barreca et al., 2016), but accounting for adaptation costs has remained a persistent challenge because they are not directly observed (Carleton and Hsiang, 2016). The costs of any adaptive measures that require resources (e.g. capital) are the opportunity costs of allocating those resources to other economic opportunities. Systematically enumerating these costs may be challenging or impossible, since there is a vast number of adjustments populations make when adapting to their climate. We point out here that in general equilibrium, all of these adaptations can simply be thought of as allocation decisions made conditional on the climate, and their total opportunity cost to the market is therefore the loss of total revenue relative to an alternative allocation. Thus, it is not necessary to enumerate these costs individually to capture their total, so long as we observe changes in total revenue in response to climate.

We begin by first defining what we mean by the word “climate.” Much of the prior economics literature considers this definition to be self-evident, but numerous debates have stemmed from disagreements about what this language represents mathematically. We

therefore provide a precise definition for climate. Our definition is novel in its articulation but we believe it both reflects and subsumes most prior conceptions and is consistent with colloquial understanding of the term. It also leads naturally to a corresponding definition of the term “weather,” reflecting the notion that “the climate is what you expect, the weather is what you get.”² These mathematical definitions are essential to our approach because they guarantee that climate and weather have the same dimensionality, lie in the same sub-regions of a vector-space, and can be mapped to data.

2.1 Defining climate and its marginal product

The climate of a location describes the joint probability distribution over a large number of possible environmental conditions that may occur at that location over a period of time. Because this distribution is a function, the transformation of climate and other inputs into economic output (a scalar) must be mapped through a *functional*.³ The challenge of our analysis is to find a suitable framework for transforming this functional into an empirically tractable object while also accounting for endogenous adaptation to any changes in climate.

Let relevant environmental state variables at location i , observed continuously in time τ , be represented by the random vector \mathbf{x} :

$$\mathbf{x}_{i\tau} = [\textit{temperature}_{i\tau}, \textit{precipitation}_{i\tau}, \textit{humidity}_{i\tau}, \dots], \quad (1)$$

where $\mathbf{x}_{i\tau}$ is a draw from the joint probability distribution function $f_{\mathbf{x}}(\cdot)$. We are interested in how changes to this function alter the economic value of output in a given economy. Assume a functional that maps this function to scalar economic output \tilde{Y} exists:

$$f_{\mathbf{x}}(\cdot) \mapsto \tilde{Y}(f_{\mathbf{x}}(\cdot), \mathbf{b}), \quad (2)$$

where \mathbf{b} is a vector of length N that describes all endogenous control variables in the economy, including the allocation of all inputs not described by \mathbf{x} . Alteration of these control variables is how the economy may adapt to changes in $f_{\mathbf{x}}(\cdot)$. Define Ψ to be the function space of physically valid probability distribution functions over \mathbf{x} , i.e. $f_{\mathbf{x}}(\cdot) \in \Psi$.

To make the characterization of (2) empirically tractable, we exploit a smooth “function-generating function” $\psi(\cdot)$ that describes the entire function⁴ $f_{\mathbf{x}}(\cdot)$ in terms of a vector \mathbf{C} :

$$\mathbf{C} \mapsto \psi(\mathbf{C}) = f_{\mathbf{x}}(\cdot). \quad (3)$$

²This quote is has been attributed to M. Twain and R. Heinlein, although its true origins are unclear.

³A functional is similar to a function, but takes functions (rather than scalars or vectors) as arguments and outputs a scalar. A definite integral is a commonly used functional.

⁴Note that $\psi(\cdot)$ may be set-valued such that it generates a set of functions $f_{\mathbf{x}}(\cdot)$ for each vector \mathbf{C} . This is not a concern for us here because each function $f_{\mathbf{x}}(\cdot)$ will still have a well-defined inverse \mathbf{C} .

Output from $\psi(\cdot)$ must always be a probability distribution. $\psi(\cdot)$ must also be sufficiently structured such that it reconstructs all economically relevant features of $f_{\mathbf{x}}(\cdot)$ based on \mathbf{C} . For example, if physics constrained Ψ to be the family of univariate Normal distributions over x , then \mathbf{C} would simply contain a mean and variance and $\psi(\mathbf{C})$ would translate these values into a function over x .

Let the vector \mathbf{C} have length K and lie in the space $\mathcal{C} \subseteq \mathbb{R}^K$. \mathcal{C} is defined such that $\mathbf{C} \in \mathcal{C} \iff \psi(\mathbf{C}) \in \Psi$. That is, for any position \mathbf{C} in the space \mathcal{C} , there exists a valid probability distribution over \mathbf{x} generated by $\psi(\mathbf{C})$. Thus, we can rewrite $\tilde{Y}(f_{\mathbf{x}}(\cdot), \mathbf{b}) = \tilde{Y}(\psi(\mathbf{C}), \mathbf{b}) = Y(\mathbf{C}, \mathbf{b})$ where the notation $Y(\cdot)$ simply embeds $\psi(\cdot)$ in $\tilde{Y}(\cdot)$. The purpose of $\psi(\cdot)$ in this formulation is that its inverse maps from the function space Ψ , which is difficult to handle empirically as the domain of an independent variable, to the vector space \mathcal{C} . This allows us to substitute the functional $\tilde{Y}(\cdot)$ with the more manageable function $Y(\cdot)$.

By construction, the vector \mathbf{C} is sufficient to describe all economically relevant features of $f_{\mathbf{x}}(\cdot)$. We therefore define \mathbf{C} as the *climate*. Correspondingly, \mathcal{C} is the *space of possible climates*. Our objective is to define and measure the economic value of relocating an economy within this space. To describe the value of such relocations, we define the *marginal product of climate* evaluated at location \mathbf{C}_0 in \mathcal{C} to be

$$\frac{dY(\mathbf{C}_0, \mathbf{b}_0^*)}{d\mathbf{C}} = \lim_{\mathbf{C}' \rightarrow 0} \left[\frac{Y(\mathbf{C}_0 + \mathbf{C}', \mathbf{b}_0^* + \mathbf{b}^{*'}) - Y(\mathbf{C}_0, \mathbf{b}_0^*)}{\mathbf{C}'} \right], \quad (4)$$

where the vectors \mathbf{b}_0^* and $\mathbf{b}_0^* + \mathbf{b}^{*'}$ are endogenously adjusted to adapt the economy to the two climates \mathbf{C}_0 and $\mathbf{C}_0 + \mathbf{C}'$, respectively. We next discuss such adaptations in more detail.

2.2 Adaptation to climate in general equilibrium

Assume a competitive market equilibrium with complete information where \mathbf{C} affects endowments, productivities, and utility (Arrow and Debreu, 1954). Firms rent labor and sell output to rational consumers according to a price-vector $\mathbf{p} = \mathbf{p}(\mathbf{C})$, which may depend on the climate. All agents know their climate \mathbf{C} and adjust their behavior and factor allocations to it independently. As mentioned earlier, the vector \mathbf{b} describes all control variables available to decision-makers, including production, consumption, and investment. Aggregate utility $U(\mathbf{C}, \mathbf{b})$ then depends on \mathbf{b} and may also depend on \mathbf{C} . Individual firms are price-takers and unable to alter $U(\mathbf{C}, \mathbf{b})$. They observe \mathbf{C} and $\mathbf{p}(\mathbf{C})$, then produce output in a decentralized manner until some equilibrium \mathbf{b}^* is achieved.

Denote the economy's production possibility frontier as $PPF(\mathbf{C}, \mathbf{b})$, which directly depends on \mathbf{C} if the climate affects what output is feasible and indirectly depends on \mathbf{C} through adaptations via \mathbf{b} . The equilibrium allocation, which incorporates all endogenous adjustments to the climate, is then $\mathbf{b}^* = \mathbf{b}^*(\mathbf{C})$. To ensure that \mathbf{b}^* exists, we assume that \mathbf{b} lies in a space \mathcal{B} , a dense subset of \mathbb{R}^N . Further, we make the usual assumption that Y

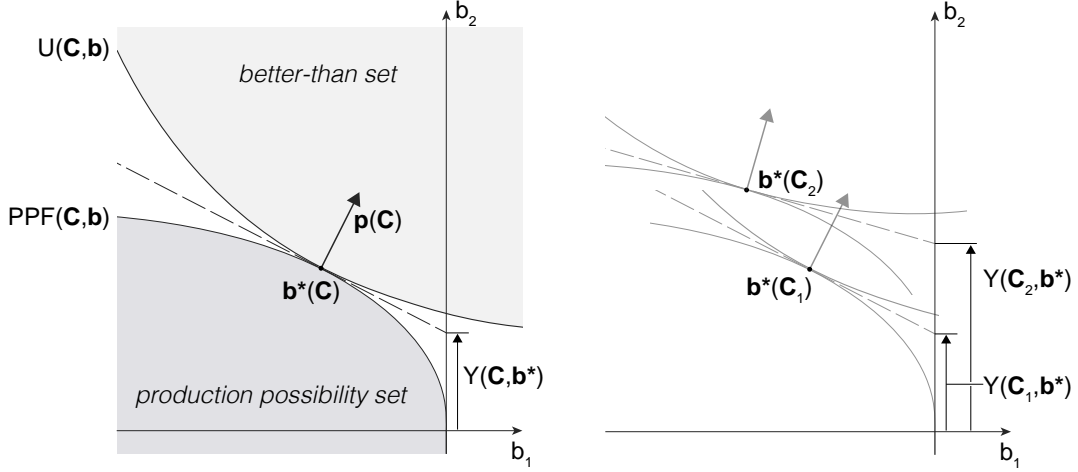


Figure 1: General equilibrium in a climate \mathbf{C} . Left: the equilibrium allocation is $\mathbf{b}^*(\mathbf{C})$ with endogenous price vector \mathbf{p} and total market revenue (income) Y . Right: The climate of an economy is altered from \mathbf{C}_1 to \mathbf{C}_2 , producing a new optimal allocation $\mathbf{b}^*(\mathbf{C}_2)$ and changing both the price vector and total revenue Y .

is concave and differentiable in \mathbf{b} . Under the conditions stated above, Koopmans (1957) shows that the equilibrium allocation $\mathbf{b}^*(\mathbf{C})$ must maximize total market revenue in the economy if it also maximizes aggregate utility, so long as prices are nonzero:

$$\mathbf{b}^*(\mathbf{C}) = \arg \max_{\mathbf{b}} Y(\mathbf{C}, \mathbf{b}(\mathbf{C})) \mid \mathbf{p}(\mathbf{C}), U(\mathbf{C}, \mathbf{b}). \quad (5)$$

Thus, as far as market revenue is concerned, all costs of adaptation to climate are opportunity costs due to the re-allocation of resources in response to changes in \mathbf{C} . If an exogenous change in the climate induces an endogenous adjustment in \mathbf{b}^* , then it must be the case that any reductions in Y caused by this adjustment were outweighed by the gains.

Figure 1 illustrates an example economy where $\mathbf{b} = (b_1, b_2)$. Both b_1 (an input to production as drawn) and b_2 (an output and the numeraire) are endogenously determined. $U(\mathbf{C}, \mathbf{b})$ and $PPF(\mathbf{C}, \mathbf{b})$ may be influenced by climate and total market revenue $Y(\mathbf{C}, \mathbf{b}^*)$ is maximized given prices (left panel). If the climate \mathbf{C} changes from \mathbf{C}_1 to \mathbf{C}_2 (right panel), both the production possibility set and the better-than set may change. The economy will adapt to \mathbf{C}_2 by reallocating factors from $\mathbf{b}^*(\mathbf{C}_1)$ to $\mathbf{b}^*(\mathbf{C}_2)$. This reallocation will again maximize total revenue Y in the economy given the new price vector. The marginal product of climate thus implicitly captures all of these adaptive factor reallocations, but measuring it does not require that we observe \mathbf{b}^* or \mathbf{p} directly. Rather, we simply observe how total revenue Y —the intercept of the separating hyper-plane with the numeraire’s axis (b_2)—responds to changes in \mathbf{C} net of these endogenous changes (right panel).

This formulation of adaptation is general in the sense that it accounts for an arbitrary number of endogenous adjustments to climate while allowing for the possibility that \mathbf{C} affects the economy by altering the better-than set and the PPF simultaneously. For example, on the production side, increasing temperatures might reduce the yields of some crops, altering the structure of the PPF and causing adjustment in the allocation of land to different crop varieties or the more intensive use of irrigation, both of which are accompanied by the opportunity cost of using those resources for other productive activities. On the preferences side, increasing temperatures might increase the demand for ice cream and reduce the demand for hot chocolate, a response that would lead to some reallocation of both production and consumption.

If some factors have no market price, such as a completely externalized pollutant or a non-market good, they will not be factored into market-based decisions and will not affect total revenue. This is not an immediate concern for us because we are focused on changes to total market output, although it is worth noting explicitly that our approach does not capture non-market responses to climate and does not necessarily correspond to welfare.⁵

Because our focus is the value of the climate in terms of total market production net of all endogenous adjustments, it is helpful to write equilibrium output as the value function

$$V(\mathbf{C}) = Y(\mathbf{C}, \mathbf{b}^*(\mathbf{C})), \quad (6)$$

which captures the net costs and benefits of all possible adaptations embodied by \mathbf{b}^* . If there exist regions in the space \mathcal{C} where endogenous adjustment of \mathbf{b} can fully offset any changes in total production induced by changes in \mathbf{C} , then the value function $V(\mathbf{C})$ will be flat in that region. However, if compensating adjustments in \mathbf{b} are not cost-effective or unfeasible, then $V(\mathbf{C})$ may have a steeper gradient that reflects the direct effect of \mathbf{C} on Y .

An important assumption above is that the elements of \mathbf{b} can take on continuous values. This is a natural assumption for most quantities in an economy of sufficiently large scale, such as the number of apples sold in a regional market or the miles of road laid down within a county. Although, in some cases often discussed as margins of adaptation to climate, such as the construction of sea-walls or crop-switching, these decisions are often framed as discrete. However, such a framing oversimplifies adaptive decisions by actual individuals because these decisions have (continuous) intensive margins. For example, a sea-wall can always be slightly longer or slightly higher and a farmer can always allocate a slightly larger fraction of cropland to a new variety. At larger scales of aggregation, such as a county, the assumption that elements of \mathbf{b} takes on continuous values is strongly defensible, as many discrete economic decisions of many individuals are aggregated into continuous measures.

⁵Examples of two alternative approaches that take into account both market and non-market responses to climate include Anthoff, Hepburn, and Tol (2009) and Hsiang et al. (2017).

2.3 Weather and climate

Once agents use their knowledge of \mathbf{C} to adjust \mathbf{b} , there remain no possible channels for the probability distribution characterized by \mathbf{C} to affect the economy except through realizations of $\mathbf{x}_{i\tau}$. A climate that is generally wet will generate more rainy minutes and a climate that is generally warm will generate more hot minutes. These actual events may affect Y , with effects possibly mediated by $\mathbf{b}^*(\mathbf{C})$. We now describe the connection between $\mathbf{x}_{i\tau}$ and what people colloquially refer to as “weather.” It is tempting to label $\mathbf{x}_{i\tau}$ the weather, but elements in $\mathbf{x}_{i\tau}$ are continuous measures taken at continuous moments in time indexed by τ . By contrast, the term “weather” typically summarizes values of $\mathbf{x}_{i\tau}$ observed over a longer time period, such as a day or month. We formalize this idea below.

Having observed some realizations of $\mathbf{x}_{i\tau}$ over the time interval $t = [\underline{\tau}, \bar{\tau})$, we can construct an empirical cumulative distribution function $\hat{F}_{\mathbf{x}}(\cdot)_{it}$ over these observations. Differentiating $\hat{F}_{\mathbf{x}}(\cdot)_{it}$ gives us $\hat{f}_{\mathbf{x}}(\cdot)_{it}$, an empirical analog to the probability density $f_{\mathbf{x}}(\cdot)$, which also must lie in the space Ψ . Noting that the function-generating function $\psi(\cdot)$ can also generate $\hat{f}_{\mathbf{x}}(\cdot)_{it}$, we define \mathbf{c}_{it} as an empirical analog to \mathbf{C} for location i during interval t :

$$\mathbf{c}_{it} \mapsto \psi(\mathbf{c}_{it}) = \hat{f}_{\mathbf{x}}(\cdot)_{it}. \quad (7)$$

By construction, \mathbf{c}_{it} has the same dimensionality as \mathbf{C} and also lies in \mathcal{C} . Put simply, \mathbf{c}_{it} summarizes the empirical distribution of many measures of $\mathbf{x}_{i\tau}$ taken over a finite interval of time. We define \mathbf{c}_{it} as the *weather* realized at i during period t .

We argue that this formal definition of \mathbf{c}_{it} maps very closely to common usages of the term “weather.” For example, if one were to ask “What was today’s weather?”, nobody would reply by describing $\mathbf{x}_{i\tau}$ directly—doing so would involve reporting a massive vector of temperatures and other variables experienced during every moment in the day. A more natural response would be to summarize all that information by describing its distribution as “pretty warm,” or “a high of 80°F and low of 60°F.” Such summary statements correspond to the elements in \mathbf{c}_{it} . After being told such summaries, it is then natural for individuals to reconstruct in their mind what a day might actually look like, i.e. the distribution of $\mathbf{x}_{i\tau}$ ’s they might experience on that day. This reconstruction procedure, taking place in the mind of an economic agent, is the transformation described by $\psi(\cdot)$.

Because \mathbf{C} summarizes the probability distribution function $f_{\mathbf{x}}(\cdot)$, which produces realizations $\mathbf{x}_{i\tau}$ that are in turn used to construct \mathbf{c}_{it} , it is straightforward to reformulate \mathbf{c}_{it} ’s as random vectors generated by some function of \mathbf{C} . Stated another way, weather is a random realization of events that are determined by the climate. To capture this intuition and simplify notation, below we refer to weather realizations \mathbf{c} determined by \mathbf{C} as $\mathbf{c}(\mathbf{C})$.

2.4 The marginal product of climate as the gradient of the value function

As suggested above, the probability distribution described by the climate \mathbf{C} can only affect economic production in two ways: (1) as information, through its effect on beliefs and subsequent economic decisions embodied by \mathbf{b}^* , and (2) through its influence on weather realizations \mathbf{c} , which in turn directly affect economic outcomes. Following Hsiang (2016), we term these pathways the *belief effect* and the *direct effect*, respectively, and we rewrite Eq. 6 to explicitly acknowledge that $\mathbf{c}(\cdot)$ is the only pathway of influence other than $\mathbf{b}(\cdot)$:

$$V(\mathbf{C}) = Y(\mathbf{c}(\mathbf{C}), \mathbf{b}^*(\mathbf{C})). \quad (8)$$

The shape of this income-generating function $Y(\cdot)$ along the optimized path $\mathbf{b}^*(\mathbf{C})$ over all $\mathbf{C} \in \mathcal{C}$ ultimately determines the value of climate $V(\mathbf{C})$. Thus, the full marginal product of the climate at \mathbf{C} can be rewritten from Eq. 4 as the local gradient in the value function

$$\frac{dY(\mathbf{c}(\mathbf{C}), \mathbf{b}^*(\mathbf{C}))}{d\mathbf{C}} = \frac{dV(\mathbf{C})}{d\mathbf{C}} = \lim_{\mathbf{C}' \rightarrow 0} \left[\frac{V(\mathbf{C} + \mathbf{C}') - V(\mathbf{C})}{\mathbf{C}'} \right], \quad (9)$$

where \mathbf{C}' describes the structure of an arbitrary perturbation to the current climate vector \mathbf{C} such that $\mathbf{C} + \mathbf{C}' \in \mathcal{C}$. Recovering the gradient vector $\frac{dV(\mathbf{C})}{d\mathbf{C}}$ for counties in the modern United States is the goal of our empirical analysis below.

2.5 Relationship to prior work

Our analysis generalizes and bridges several previous innovative efforts to empirically measure the economic impact of climatic conditions (a complete discussion is in the Appendix). Cross-sectional analyses of farms by Mendelsohn, Nordhaus, and Shaw (1994) and Schlenker, Hanemann, and Fisher (2006) estimated analogs of Equation 8 with \mathbf{C} being captured by long-run averages, the former using simple averages of temperatures and rainfall and the latter utilizing degree-days. Deschênes and Greenstone (2007) raised the concern that unobservable heterogeneity across farms might introduce bias to these estimate, and proposed differencing out these effects using a within-unit panel regression approach. However, this strategy required the assumption that climate can be proxied by observed average weather ($\hat{\mathbf{C}} = \mathbf{c}_{it}$). Such assumptions have subsequently been widely challenged because economic agents adapt to their climate, adjusting $\mathbf{b}^*(\mathbf{C})$, but do not adapt to weather. Later work explored the nature of such adaptation in multiple ways: Schlenker and Roberts (2009); Aroonruengsawat and Auffhammer (2011); Hsiang and Narita (2012) and others showed that historical experience with climates altered the marginal impact of weather; Deschênes and Greenstone (2011); Barreca et al. (2016) and others showed how specific technologies have mitigated effects of weather; Dell, Jones, and Olken (2012); Burke and Emerick (2016) and others showed that the scope for adaptation may sometimes be limited even over longer

periods; and Schlenker, Roberts, and Lobell (2013); Houser et al. (2015) and others point to the importance of accounting for both costs and benefits of adaptation.

Thus, the literature is overall in broad agreement that (i) climate and weather are different objects, (ii) long-run climates are correlated with unobservable heterogeneity that hinders direct empirical inference in cross-sections, (iii) adaption to climate may generate important differences between the effects of climate and weather, and (iv) the costs and benefits of adaptation must both be accounted for in valuations. Yet to date, no formulation has accounted for all of these concerns in a single empirical framework. Below we show how the market value of climate—purged of bias from unobservable heterogeneity and accounting for both the costs and benefits of adaptation—can be empirically recovered by exploiting weather variation and a within-unit panel estimator.

3 Identifying the full marginal product of climate empirically

As formulated here, an economy’s output Y depends on its position in the $K+N$ dimensional space $\mathcal{C} \times \mathcal{B}$. Exogenous changes in the position of the economy $\mathbf{C} \in \mathcal{C}$ lead to endogenous re-optimization of control variables $\mathbf{b} \in \mathcal{B}$ such that, in the long run, income $Y(\mathbf{c}(\mathbf{C}), \mathbf{b}^*(\mathbf{C})) = V(\mathbf{C})$ is at an optimum with respect to the subspace \mathcal{B} . A core empirical challenge has been tracing out the path of an economy through $\mathcal{C} \times \mathcal{B}$ as it adjusts \mathbf{b}^* in response to changes in \mathbf{C} . To that end, our objective is to characterize the marginal product of climate⁶ along this path:

$$\frac{dV(\mathbf{C})}{d\mathbf{C}} = \left[\frac{\partial V(\mathbf{C})}{\partial \mathbf{C}_1}, \dots, \frac{\partial V(\mathbf{C})}{\partial \mathbf{C}_K} \right], \quad (10)$$

where \mathbf{C}_k denotes the k -th element in \mathbf{C} . To characterize $\frac{dV(\mathbf{C})}{d\mathbf{C}}$ empirically, we decompose it into contributions that come from (1) the direct effect of changing \mathbf{C} on weather realizations \mathbf{c} (direct effect) and (2) the endogenous adjustment of \mathbf{b} in response to the knowledge or belief that \mathbf{C} has changed (belief effect):

$$\frac{dV(\mathbf{C})}{d\mathbf{C}} = \frac{dY(\mathbf{c}(\mathbf{C}), \mathbf{b}^*(\mathbf{C}))}{d\mathbf{C}} = \underbrace{\sum_{k=1}^K \frac{\partial Y}{\partial \mathbf{c}_k} \frac{d\mathbf{c}_k}{d\mathbf{C}}}_{\text{direct effects}} + \underbrace{\sum_{n=1}^N \frac{\partial Y}{\partial \mathbf{b}_n} \frac{d\mathbf{b}_n^*}{d\mathbf{C}}}_{\text{belief effects}}, \quad (11)$$

where $\frac{d\mathbf{c}_k}{d\mathbf{C}}$ and $\frac{d\mathbf{b}_n}{d\mathbf{C}}$ are the k -th and n -th row vectors of Jacobians $\frac{d\mathbf{c}}{d\mathbf{C}}$ (size $K \times K$) and $\frac{d\mathbf{b}}{d\mathbf{C}}$ (size $N \times K$), respectively.⁷ Importantly, each partial derivative is evaluated “locally” in the neighborhood of the initial climate \mathbf{C} and its associated equilibrium $\mathbf{b}^*(\mathbf{C})$.

⁶Note that the marginal product of climate is equivalent to the gradient of the value function $\nabla V(\mathbf{C})$, which always points in the direction of most rapid ascent, locally.

⁷ The Jacobian matrices are $\frac{d\mathbf{c}}{d\mathbf{C}} = \begin{bmatrix} \frac{\partial \mathbf{c}_1}{\partial \mathbf{C}_1} & \dots & \frac{\partial \mathbf{c}_1}{\partial \mathbf{C}_K} \\ \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{c}_K}{\partial \mathbf{C}_1} & \dots & \frac{\partial \mathbf{c}_K}{\partial \mathbf{C}_K} \end{bmatrix}$ and $\frac{d\mathbf{b}}{d\mathbf{C}} = \begin{bmatrix} \frac{\partial \mathbf{b}_1}{\partial \mathbf{C}_1} & \dots & \frac{\partial \mathbf{b}_1}{\partial \mathbf{C}_K} \\ \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{b}_N}{\partial \mathbf{C}_1} & \dots & \frac{\partial \mathbf{b}_N}{\partial \mathbf{C}_K} \end{bmatrix}$.

In the ideal experiment, we would compare two identical economies with initial climate \mathbf{C} , perturbing the climate of one by a small disturbance \mathbf{C}' to $\mathbf{C} + \mathbf{C}'$ and allowing the economy to adjust \mathbf{b}^* endogenously. The resulting difference in productivities would then be the marginal product $\frac{dV(\mathbf{C})}{d\mathbf{C}}$, capturing contributions from both direct effects and belief effects. As highlighted in the literature, the central empirical challenge to finding suitable quasi-experimental conditions to measure $\frac{dV(\mathbf{C})}{d\mathbf{C}}$ empirically has been the absence of reliable exogenous changes in \mathbf{C} . For example, long-run average climate conditions may be correlated with unobserved heterogeneity (Deschênes and Greenstone, 2007); and while gradual changes of the climate within a location may be available in some contexts, they may be correlated with unobserved trends in confounding variables that also evolve slowly (e.g. the *frequency-identification tradeoff* described in Hsiang and Burke, 2014). Thus, it is appealing to examine random fluctuations in weather \mathbf{c} as proposed by Deschênes and Greenstone (2007), since this source of variation is generally orthogonal to unobserved heterogeneity. However, it has remained an open question whether $\frac{dY}{d\mathbf{c}}$ can be mapped to $\frac{dV}{d\mathbf{C}}$ in a manner consistent with economic theory because populations adapt to climate (via \mathbf{b}) but not to weather and because the belief effects in Eq. 11 are likely too numerous to characterize through enumeration. A key insight of this paper is that, while these two gradients are different mathematical objects, their values *are* equal in the neighborhood of \mathbf{b}^* , allowing us to measure the value of climate by isolating the marginal product of weather net of unobserved heterogeneity. This insight results from application of the Envelope Theorem.⁸

The intuition is as follows. Imagine there are identical adjacent villages along a road that runs North-South, and each village has a single choice variable, i.e., $\mathbf{b} = b$. Each village also faces a one-dimensional climate $\mathbf{C} = C$, with more northern villages experiencing colder conditions. The surface $Y(C, b)$ is illustrated in the left panel of Figure 2.⁹ Given their climate C , individuals and firms in the villages maximize their utility and profit, respectively. The result is that each village’s b^* is such that Y is maximized by firms given prices. Therefore, in equilibrium, villages’ outputs lie along the “ridge” of the Y surface (viewed from the perspective of firms), which is equal to the value function $V(C)$ indicated by the blue line. An observer traveling along this road would observe that village economies differ in b , since villages adapt as C changes. In contrast, if villages did not adapt, b would be constant regardless of C , locating output along the road on a “slice” of the Y surface at a fixed b rather than the ridge. Two examples of such slices are depicted as red and orange lines for b_1 and b_2 , respectively, each of which is optimal for a single village located where

⁸In related work, Guo and Costello (2013) exploit the Envelope Theorem to demonstrate that adaptation to climate should generate limited value on the margin in California timberland management. Similarly, Schlenker, Roberts, and Lobell (2013) demonstrate empirically that marginal costs of adaptation to temperature in US maize production closely match marginal benefits at the current equilibrium, a result fully consistent with the predictions of the Envelope Theorem as it is used in the present analysis.

⁹In the Appendix A2, we expand on this example for the case with a two-dimensional climate.




Figure 2: Output in a one-dimensional climate C with one control variable b . Left: Gray surface is output $Y(C, b)$, blue line is the value function $Y(C, b^*(C)) = V(C)$, red and orange lines are “slices” where b is held at fixed values. Right: Displaying $V(C)$ and slices from the left panel when they are projected onto the $Y-C$ plane. Appendix Section A2 displays the value function for a two-dimensional climate.

the slice is crossed by the blue line.

If we do not observe b for each village, then we only see the relationships between Y and C , shown as the projection of the equilibrium (blue) onto the $Y-C$ plane on the right panel of Figure 2. The slices for $b = b_1$ and $b = b_2$ are also projected onto this plane, and lie below the equilibrium everywhere except at the climate for which they are optimal.

Now, if the weather of a village changes unexpectedly such that it is different from its climate by ξ , b^* will not be adjusted in the short run and the village will produce $Y(C + \xi, b^*(C))$. Output of the village will be determined by its respective slice of Y where b is held fixed at $b^*(C)$ (e.g. the red line). For large perturbations, output will temporarily fall below the ridge (blue line), but so long as the village’s beliefs about their climate do not change, they will remain at $b^*(C)$ and output will recover when the weather returns to C . For this reason, numerous researchers have argued that econometric measurements that exploit weather variation as the independent variable do not recover the effect of altering C once b is allowed to adjust (reviewed in Dell, Jones, and Olken, 2013; Hsiang, 2016). Here we point out that, for small perturbations, the responses of Y to weather (red and orange lines) are tangent to the value function (blue line) in the neighborhood of each village’s b^* . Thus, even if the value function cannot be observed directly, its local derivative—i.e. the marginal product of climate—can be measured locally by exploiting variation in weather where ξ is small. Then, if we can observe a large number of these local derivatives for nearby villages with different initial C , we can piece them together via integration to recover the overall shape of the value function along this hypothetical road where adaptations occur continuously across neighboring villages.

Therefore, the first key result is the equivalence, locally, of the marginal effect of climate and the marginal effect of weather in the neighborhood of an initial equilibrium, regardless of any possible margins of adaptation. This equivalence is obtained because (i) no adaptation occurs in response to weather and (ii) adaptation has no influence on the marginal effect of climate locally since Y is initially optimized. That is, the effect of adaptation locally is zero in both cases. The second key result is that a large number of local marginal effects can be integrated to compute the non-marginal effect of large changes in the climate.

To see these results in the general case, let \mathbf{C}^a be a benchmark climate at which we are evaluating the marginal product of climate. To estimate the k -th element of $\frac{dV(\mathbf{C}^a)}{d\mathbf{C}}$ (Eq 10), we differentiate V by \mathbf{C}_k at \mathbf{C}^a . By the chain rule, we have

$$\frac{dV(\mathbf{C}^a)}{d\mathbf{C}_k} = \underbrace{\frac{\partial Y(\mathbf{c}(\mathbf{C}^a), \mathbf{b}^*(\mathbf{C}^a))}{\partial \mathbf{C}_k}}_{=0} + \underbrace{\sum_{\kappa=1}^K \frac{\partial Y(\mathbf{c}(\mathbf{C}^a), \mathbf{b}^*(\mathbf{C}^a))}{\partial \mathbf{c}_\kappa} \frac{d\mathbf{c}_\kappa}{d\mathbf{C}_k}}_{\text{direct effects}} + \underbrace{\sum_{n=1}^N \frac{\partial Y(\mathbf{c}(\mathbf{C}^a), \mathbf{b}^*(\mathbf{C}^a))}{\partial \mathbf{b}_n} \frac{d\mathbf{b}_n}{d\mathbf{C}_k}}_{\sum \text{belief effects}=0}. \quad (12)$$

The climate, as a probability distribution, cannot affect any outcome by a pathway other than through the weather realizations it causes and actions based on beliefs regarding its structure. This implies the first term must be zero. Because V is the outcome when Y has been optimized through each possible adaptation \mathbf{b}_n , we also know that $\frac{\partial Y}{\partial \mathbf{b}_n} = 0$ for all n . This means the sum of all belief effects must be zero *even if observable adaptive adjustments* $\frac{d\mathbf{b}}{d\mathbf{C}_k}$ *are large*. Thus, only the terms $\frac{\partial Y}{\partial \mathbf{c}_\kappa} \frac{d\mathbf{c}_\kappa}{d\mathbf{C}_k}$ (direct effects) may be nonzero.

Next, we note that for any marginal change in the distribution of weather, there exists a marginal change in climate that is equal in magnitude and structure such that $\frac{d\mathbf{c}_\kappa}{d\mathbf{C}_k} = 1$ if $\kappa = k$ and 0 otherwise. Thus, we can focus only on cases where $\kappa = k$, i.e. the effect of the κ -th element of \mathbf{c} is thought to be informative of the effect of the k -th element of \mathbf{C} .¹⁰ Then we have $\frac{dV(\mathbf{C}^a)}{d\mathbf{C}_k} = \frac{\partial Y(\mathbf{C}^a, \mathbf{b}^*)}{\partial \mathbf{c}_k}$, which says that the total marginal effect on V of the k th dimension of the climate, evaluated at \mathbf{C}^a , is equal to the partial derivative of income with respect to the corresponding dimension of weather, also evaluated at \mathbf{C}^a . Locally, the marginal effect of the climate on income is identical to the marginal effect of the weather. Extending this to all K dimensions of the climate we have our first result

$$\frac{dV(\mathbf{C}^a)}{d\mathbf{C}} = \frac{\partial Y(\mathbf{C}^a, \mathbf{b}^*)}{\partial \mathbf{c}}, \quad (13)$$

stating that the full K -dimensional marginal product of climate, net of all endogenous adaptations, is equal to the vector of partial effects of K weather measures, ignoring all

¹⁰ This restriction is equivalent to setting $\frac{d\mathbf{c}}{d\mathbf{C}}$ equal to the identity matrix and is quite weak. It simply requires that we do not interpret changes in one measure of weather (e.g. realized average temperature) as reflecting changes in an orthogonal climate measure (e.g. expected rainfall).

adaptations.¹¹ Equation 13 is particularly useful empirically because the right-hand-side term can be estimated in a multivariate time-series or panel model regression that is purged of location-specific heterogeneity (e.g. using county fixed effects), following the approach laid out in Deschênes and Greenstone (2007). Importantly, however, for Equation 13 to hold, the outcome must represent a maximized quantity, which is true in the case of income—studied here—but may not hold for other outcomes, such as crop yields or mortality risk.

Empirical estimates of Equation 13 can be used to construct estimates of non-marginal climate effects by integrating marginal effects of weather. For an arbitrary climate \mathbf{C}^b , we can solve for $V(\mathbf{C}^b)$ by computing a line integral of the gradient in V along a continuous path through the K -dimensional \mathbf{C} -subspace (i.e. \mathcal{C}) from $\mathbf{C}^a \rightarrow \mathbf{C}^b$, starting from $V(\mathbf{C}^a)$:

$$V(\mathbf{C}^b) = V(\mathbf{C}^a) + \int_{\mathbf{C}^a}^{\mathbf{C}^b} \frac{dV(\mathbf{C})}{d\mathbf{C}} d\mathbf{C}. \quad (14)$$

At each position $\mathbf{C} \in \mathcal{C}$, $\frac{dV(\mathbf{C})}{d\mathbf{C}}$ is a vector of differentials describing all the marginal effects of the climate measured locally at \mathbf{C} . From Equation 13 we know that these differentials with respect to climate can be substituted for using differentials with respect to weather, which in turn can be estimated empirically via regression. Empirically, $\frac{\partial Y(\mathbf{C}, \mathbf{b}^*)}{\partial \mathbf{c}} = \hat{\beta}_{weather} \Big|_{\mathbf{C}}$, where $\hat{\beta}_{weather}$ is a reduced-form parameter estimate of the marginal effects of weather on total economic production. Combining this fact with Eq. 13, we obtain our second result, an estimate for the net value of a non-marginal change in the climate from \mathbf{C}^a to \mathbf{C}^b exploiting only exogenous variation in weather:

$$V(\mathbf{C}^b) - V(\mathbf{C}^a) = \int_{\mathbf{C}^a}^{\mathbf{C}^b} \hat{\beta}_{weather} \Big|_{\mathbf{C}} \cdot d\mathbf{C}. \quad (15)$$

Eq. 15 says that the change in value resulting from a non-marginal change in climate can be computed as a line integral through a vector field of empirical gradient estimates constructed by regressing income on idiosyncratic weather variations.

Because non-marginal changes in climate may induce substantial adjustments to control variables in the economy, $\frac{dV(\mathbf{C})}{d\mathbf{C}}$ may change with \mathbf{C} , implying that fully accounting for adaptation in Eq. 15 requires that new marginal effects $\hat{\beta}_{weather}$ be empirically estimated at each position in \mathcal{C} . If these marginal effects are not constant across \mathcal{C} , then $V(\mathbf{C})$ has curvature in the K -dimensional space. If a single marginal effect of weather is estimated, pooling across many baseline climates, this is equivalent to forcing the marginal effect of climate to be constant in each dimension of \mathcal{C} . Such a “constant marginal effect model” might reasonably approximate $V(\mathbf{C})$, although it is difficult to be certain whether this is

¹¹Hsiang (2016) describes the assumption of Eq. 13 in prior work as the *marginal treatment comparability assumption*, which holds exactly in this context.

the case *ex ante*. In our empirical implementation, we explore when constant marginal effects are a good approximation and when nonlinearity is important for capturing adaptive responses to climate expressed as curvature in $V(\mathbf{C})$.

4 Empirical implementation for the modern United States

In our empirical analysis of the temperature climate for the US, we define $\psi(\cdot)$ such that the probability distribution of daily average temperatures within each year and location is described by a 17-bin histogram, where the interior fifteen bins are each 3°C wide and the top and bottom bins are not bounded above and below, respectively. The space of possible climates \mathcal{C} is thus the 16-simplex, constrained such that the total number of days in a year is exactly 365. The position of a county i on the 16-simplex \mathcal{C} is then \mathbf{C}_i , a 16-element-long vector describing the expected count of days in each temperature bin less one.

4.1 Identification strategy

Our objective is to empirically estimate the value of permanently repositioning county i 's climate \mathbf{C}_i to some location $\mathbf{C}_i + \mathbf{C}'$. To measure this value, our empirical strategy exploits the result in Equation 13 and measures the marginal value of changes in the weather vector \mathbf{c}_{it} within location i over time, which is equal to the marginal value of the permanent distortion in \mathbf{C}_i . The weather vector \mathbf{c}_{it} necessarily has the same 16-element structure as \mathbf{C}_i and can be written as $\mathbf{c}_{it} = \mathbf{C}_i + \xi_{it}$, where ξ_{it} is a vector of disturbance terms.¹² As ξ_{it} varies randomly over time, the position of \mathbf{c}_{it} will “explore” \mathcal{C} in the neighborhood of \mathbf{C}_i , allowing us to estimate the local marginal effect of these changes on output. An important econometric benefit of exploiting the within-county variation in \mathbf{c}_{it} is that it allows us to utilize panel-regression techniques that condition out unobservable heterogeneity across counties using a fixed effect (Schlenker and Roberts, 2009; Deschênes and Greenstone, 2011). The left panel of Figure 3 displays an example realization of \mathbf{c}_{it} overlaid on the climate \mathbf{C}_i . The difference between these distributions (black outlines, bottom) is the vector ξ_{it} , which provides our identifying variation.

Once we empirically estimate the local structure of $V(\mathbf{C}_i)$ by exploiting within-county variation in \mathbf{c}_{it} , we can use these local marginal effects to compute $V(\mathbf{C}_i + \mathbf{C}')$ via integration (Equation 15). The right panel of Figure 3 heuristically illustrates how we might apply this approach to compute the difference in value between the climate of St. Paul, Minnesota and Orlando, Florida. By tracing a path through adjacent counties from St. Paul to Orlando, we gradually move through the 16-simplex of \mathcal{C} and at each step use local variation in weather to estimate $\frac{\partial V(\mathbf{C}_i)}{\partial \mathbf{c}_{it}}$, which we then integrate to determine the shape of the surface from $V(\mathbf{C}_i)$ to $V(\mathbf{C}_{i+1})$. It is worth noting that the exact path taken from St. Paul to Orlando should not matter, so long as $V(\mathbf{C})$ is sufficiently smooth and counties are sufficiently “near”

¹²Note that it need not be the case that $E[\mathbf{c}_{it}] = \mathbf{C}_i$ in order to trace out a tangency to the value function.

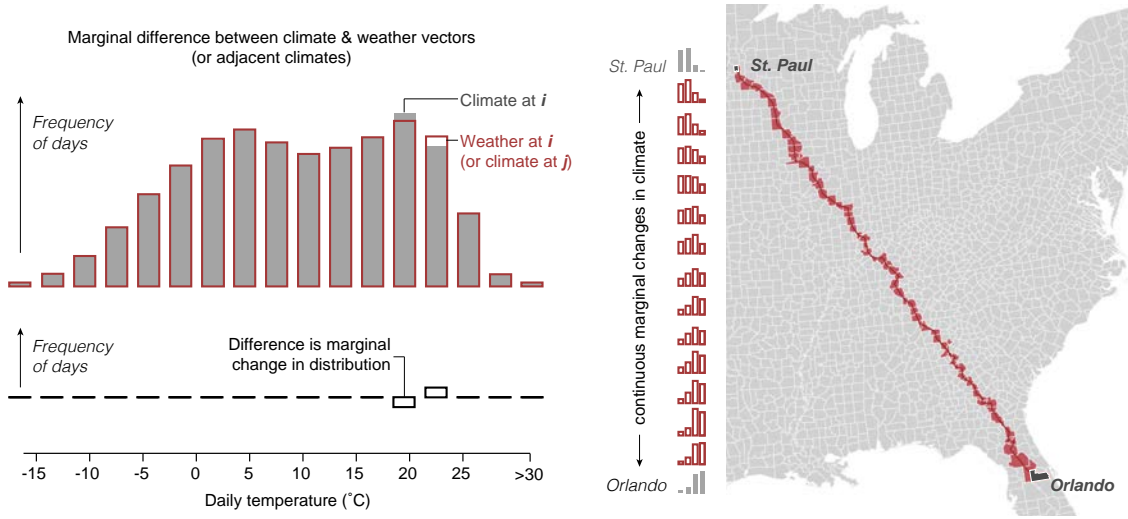


Figure 3: (Left) Comparing the average temperature distribution of a county, where bin heights are C_i , (gray) with a slightly altered distribution resulting from the weather c_{it} (red outline), which might be identical to C_j for adjacent county j . The difference between these distributions (black outline) is the identifying variation in temperature distributions (ξ_{it}). (Right) Example integration path (red) used to compute the difference in value of the climate in St. Paul, MN versus Orlando, FL (both black). Small histograms illustrate hypothetical temperature distributions along this path, evolving by almost continuous marginal changes.

one another in \mathcal{C} , relative to the curvature of $V(\mathbf{C})$. The more linear $V(\mathbf{C})$ is, the more reasonable it is to extrapolate between counties distant in \mathcal{C} .

4.2 Regression specification

We construct our empirical specification to closely reflect the theoretical structure of the climate value function. To do this, we combine Equations 10, 13, and 14, setting the benchmark climate to C_i . Additionally, we allow the overall level of the value function to depend on a variety of other unobserved location- and time-specific factors \mathbf{z}_{it} that may be unrelated to the climate, such as history and human capital endowments, which affects the county-specific time-varying constant of integration $V(C_i|\mathbf{z}_{it})$:

$$\begin{aligned}
Y(C_i + C'_{it}, \mathbf{b}^*) &= V(C_i + C'_{it}) = V(C_i|\mathbf{z}_{it}) + \int_{C_i}^{C_i+C'_{it}} \frac{dV(\mathbf{C})}{d\mathbf{c}} \cdot d\mathbf{c} \\
&= V(C_i|\mathbf{z}_{it}) + \int_{C_i}^{C_i+C'_{it}} \left[\frac{\partial V(\mathbf{C})}{\partial c_{k=1}}, \dots, \frac{\partial V(\mathbf{C})}{\partial c_{k=K}} \right] \cdot d\mathbf{c} \\
&= V(C_i|\mathbf{z}_{it}) + \left[\frac{\partial V(C_i)}{\partial c_{k=1}}, \dots, \frac{\partial V(C_i)}{\partial c_{k=K}} \right] \cdot \xi_{it}. \quad (16) \\
\mathbf{C}' &\rightarrow 0 \\
\xi_{it} &\rightarrow 0
\end{aligned}$$

The term $V(\mathbf{C}_i|\mathbf{z}_{it})$ is the total income that the economy at i would obtain at time t conditional on covariates \mathbf{z}_{it} if the climate remained at \mathbf{C}_i . In the empirical specification, we model the influence of \mathbf{z}_{it} using vectors of county fixed-effects, year fixed-effects, and an auto-regressive term. The last equality in Eq. 16 is the first-order Taylor expansion of $V(\mathbf{C}_i + \xi_{it})$ and holds exactly for sufficiently small changes in ξ . The last term in this line is the inner product between the K -dimensional gradient vector of the value function evaluated at \mathbf{C}_i (the object of interest) and the disturbance vector ξ_{it} describing how weather conditions in county i and time t deviate from \mathbf{C}_i .

Importantly, the marginal product of climate is identified locally, in the neighborhood of \mathbf{C}_i , and it is possible that marginal products at different positions in \mathcal{C} are not identical. This may occur if, for example, populations adapt to climates in ways that alter the local marginal product (e.g. installing air conditioning). To account for non-constant marginal effects, we construct a model that is nonlinear in each of the K dimensions of \mathbf{C} , which allows the marginal product of climate to change as a function of the position $\mathbf{C}_i \in \mathcal{C}$, thereby fully capturing all effects of adaptation as well the net effect of all other non-linear responses in the PPF (e.g. Schlenker and Roberts, 2009) or aggregate demand (e.g. Auffhammer, Baylis, and Hausman, 2017).

Next, we construct an empirical analog to Equation 16. The dimensions of climate we consider in our main specification include daily temperature and precipitation distributions in both the current and past year. Our main focus is on the effect of current temperatures, however, as these other dimensions of the climate appear to have little effect on the value function. We also account for within-county autocorrelation, unobserved heterogeneity across counties, and nonlinear time trends. Specifically, we estimate

$$Y_{it} = \rho Y_{i,t-1} + \mu_i + \theta_t + \sum_{h=1}^H \left[\sum_m \left[\beta^{mh} (\tilde{T}_{it}^m)^h + \gamma^{mh} (\tilde{T}_{i,t-1}^m)^h \right] \right] + \sum_g \left[\zeta^n \tilde{P}_{it}^g + \eta^n \tilde{P}_{i,t-1}^g \right] + \epsilon_{it}, \quad (17)$$

where counties are indexed by i and years are indexed by t . Y_{it} is a measure of output, which in our main specification is *log income per capita*. μ_i is a set of county fixed effects that account for unobserved constant differences between counties, such as elevation. 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. The model allows the value function to be nonlinear in each dimension (bin) of the temperature climate up to the order H , which is crucial for fully accounting for adaptation.¹³ In our “constant marginal effects model,” we set $H = 1$, thereby constraining the marginal effect of a hot day to remain fixed throughout the support of the weather data. We then re-estimate

¹³We find that the role of precipitation is essentially zero, even in an affine model, so we do not present models that account for curvature in the precipitation subspace of \mathcal{C} .

Equation 17 with $H = 3$, allowing for independent and asymmetric curvature in every temperature dimension of \mathcal{C} . This allows, for example, the marginal effect of additional hot days to become larger or smaller depending on whether a location already experiences a large number of hot days.

\tilde{T}_{it}^m is the number of days in county i and year t that have 24-hour average temperatures (T_d) falling in the m th temperature bin. 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.¹⁴ The top ($m = 17$) bin counts days with $T_d \geq 30^\circ\text{C}=86^\circ\text{F}$. The omitted category is the bin for $T_d \in [12, 15)^\circ\text{C} = [53.6, 59)^\circ\text{F}$. Daily precipitation bin values \tilde{P}^g are defined similarly: each of the 12 precipitation bins spans 40mm of rainfall equivalent, 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 across years within a county, we include lagged values for all \tilde{T}^m and \tilde{P}^g variables to capture any possible direct effects that weather in the prior year might have on current output. We explore additional lags (across time and space) in extensions.

$Y_{i,t-1}$ is a lagged dependent variable with serial correlation coefficient ρ . Including this term 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 affects production in subsequent years. One drawback of dynamic panel models, such as Equation 17, 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. The magnitude of potential bias in our case is less than 5% of the magnitude of our point estimate, far smaller than our uncertainty due to sampling error.¹⁵

Finally, 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 county.¹⁶

¹⁴For display purposes, coefficients on the two coldest temperature bins ($T_d < -15^\circ\text{C}$ and $T_d \in [-12, -9)^\circ\text{C}$) are not shown in figures. Generally, there are few observations at these extremely cold temperatures, and the estimated effects are highly uncertain and not statistically different from zero.

¹⁵Nickell (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, $\frac{-(1+\rho)}{(T-1)}$ is approximately 0.045 in our case. For completeness, we have also computed estimates without any lagged dependent variable and continue to obtain our main result (available upon request).

¹⁶See Fisher et al. (2012) for a discussion of this technique to account for spatial autocorrelation and Hsiang (2010) for a discussion of simultaneously accounting for spatial and temporal autocorrelation.

The coefficients β^{mh} are the parameters of interest, as they characterize 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 the constant marginal effects model ($H = 1$), the vector of coefficient estimates $[\hat{\beta}^{m=1}, \dots, \hat{\beta}^{m=16}]$ is directly interpretable as the gradient vector for the value function, the marginal product of climate $\widehat{\frac{dV(\mathbf{C})}{d\mathbf{C}}}$, as indicated by the inner-product term in Eq. 16. When we set $H = 3$ to fully account for adaptation, interpretation becomes more nuanced. The vector of coefficients (which is now three times longer) can no longer be directly interpreted as the gradient vector of the value function because the first-order Taylor expansion used in Equation 16 is no longer exact. Instead, the total effect of \tilde{T}_{it}^m days in the m th temperature bin is estimated as the polynomial $\hat{\beta}^{m1}\tilde{T}_{it}^m + \hat{\beta}^{m2}(\tilde{T}_{it}^m)^2 + \hat{\beta}^{m3}(\tilde{T}_{it}^m)^3$. Thus, the marginal effect of each additional day in the m th temperature bin—i.e. m th element in $\widehat{\frac{dV(\mathbf{C})}{d\mathbf{C}}}$ —is the derivative of this polynomial with respect to the count of days in the m th bin: $\hat{\beta}^{m1} + 2\hat{\beta}^{m2}\tilde{T}_{it}^m + 3\hat{\beta}^{m3}(\tilde{T}_{it}^m)^2$.

4.3 Data

We use county-level weather and income data for the lower 48 states over the period 1969-2011.¹⁷ To measure daily maximum and minimum temperatures as well as precipitation, we use surface data from the National Centers for Environmental Information (NCEI).¹⁸ We match weather stations to counties using each station’s reported latitude and longitude. We omit cases where the maximum or minimum temperature exceeds 60°C or is lower than -80°C, as these are likely errors. If there are multiple stations within a county, we average their measures for each day. Our measure of daily temperature is a simple average between the maximum and minimum temperatures, which is the standard measure for average temperature during a 24-hour period.¹⁹ We drop county-by-year observations that do not have a complete set of daily observations, as their full daily temperature distribution is unknown.²⁰

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.²¹ The most comprehensive income measure at the county level is *total personal income*. It includes 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. Wages and salaries include tips, commissions, bonuses, and any “pay-

¹⁷For additional data details, including summary statistics for key variables, see the Appendix.

¹⁸Available from ftp://ftp.ncdc.noaa.gov/pub/data/ghcn/daily/by_year/.

¹⁹Diurnal temperatures cycles are roughly sinusoidal, so this is a good approximation for the true mean.

²⁰As discussed in Auffhammer et al. (2013), weather station data is often incomplete due to mechanical failures or political events.

²¹For further details, see <http://www.bea.gov/regional/pdf/lapi2016.pdf>.

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 laws.

REIS measures of *farm income* are derived from United States Department of Agriculture estimates, which are based on sample surveys, Agricultural Census data, and administrative data. We use *cash receipts from marketing crops* to capture gross income. We also examine *net farm income*, which includes additions to inventories, transfers such as subsidies, crop insurance, and disaster payments (we examine these transfers in the Appendix). We inflation-adjust all measures to 2011 dollars and convert them to per capita terms.

5 The marginal product of temperature

5.1 Results assuming constant marginal effects

We first present estimates of the marginal product of temperature assuming the value function $V(\mathbf{C})$ is a flat hyperplane spanning a 16-dimensional \mathcal{C} defining all possible contemporaneous temperature distributions, i.e. $H = 1$ in Equation 17. This model provides a good first-order approximation of the marginal product of temperature.²² In the next section we demonstrate how allowing for curvature in the value function alters these results.

The left panel in Figure 4 shows our main result: the marginal product of daily temperature with respect to personal income per capita. Specifically, the figure displays the vector of coefficient estimates for contemporaneous daily temperatures $[\hat{\beta}^{m=1}, \dots, \hat{\beta}^{m=16}]$, which is exactly equal to the estimated marginal product of climate $\widehat{\frac{dV(\mathbf{C})}{d\mathbf{C}}}$. Because this vector is the gradient of $V(\mathbf{C})$, each point estimate $\hat{\beta}^m$ can be interpreted as the marginal effect of increasing the annual count of days in that temperature bin by one (Deschênes and Greenstone, 2011). Implicitly, such a change requires removing a day from the omitted temperature bin to ensure a valid annual temperature distribution that lies in \mathcal{C} .

We find that counties’ log personal income per capita increases slightly as temperatures rise from cool to moderate, then declines approximately linearly at 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. In the Appendix, we consider the effect of temperature on the *earnings* component of personal income and obtain similar, albeit larger, results.

For a sense of magnitudes, note that if output were uniform across 365 days in a year, then each day would contribute $\frac{1}{365} = 0.27\%$ of annual income. In this case, a decline of 0.065% of annual income from an extremely hot day with average temperature of 29°C (84.2°F) indicates a productivity loss of roughly 23.6% relative to an average day. Linearizing the effect of temperature relative to the zero effect at 15°C (59°F), this is a marginal change in daily productivity of $\frac{-23.6\%}{(29-15)^{\circ}\text{C}} = -1.68\%/^{\circ}\text{C} = -0.93\%/^{\circ}\text{F}$.

²²Our estimates of the effect of rainfall, available upon request, are essentially zero.




Figure 4: The estimated marginal product of contemporaneous (left) and lagged annual temperature distributions. Ordered point estimates plot elements of the gradient vector $\frac{dV(\mathbf{C})}{d\mathbf{C}}$. Values can be interpreted as the relative effect of distorting the annual temperature distribution by increasing the count of days at the indicated temperature by one. Outcome is log annual total personal income per capita $\times 100$ (i.e. percentage points) in US counties for 1969-2011. For reference, an average day contributes $\frac{1}{365} = 0.27\%$ of annual income. Shaded area denotes 95% confidence intervals for comparisons against the omitted category $T \in [12, 15)^\circ\text{C}$. Left and right panels are estimated simultaneously in a single regression model.

The right panel in Figure 4 displays the estimated effect of daily temperatures on annual income per capita the *following* year (γ^m in Equation 17). We estimate these effects jointly with the contemporaneous effect shown in the left panel. Except for the hottest temperature bin ($> 30^\circ\text{C}$), we do not observe any statistically significant effect of daily temperatures on income the following year. Even this single significant coefficient may be spurious, as we are testing sixteen coefficients. Alternatively, if interpreted as meaningful, this coefficient suggests that roughly half of the income loss from the very hottest days is made up in the following year.²³

Persistence A key requirement of our approach is that the weather events we exploit for identification do not perturb an economy so far from its equilibrium that the Envelope Theorem is no longer valid. So far, one indication that this assumption holds here is the finding that changing the temperature of an individual day lowers annual output by only fractions of a percentage point. A second indication that we now test for is whether the economy moves back to its initial equilibrium after an economic disturbance triggered by a single

²³In later simulations, we include temperature lags to ensure we do not mis-estimate the effect of high temperature days.

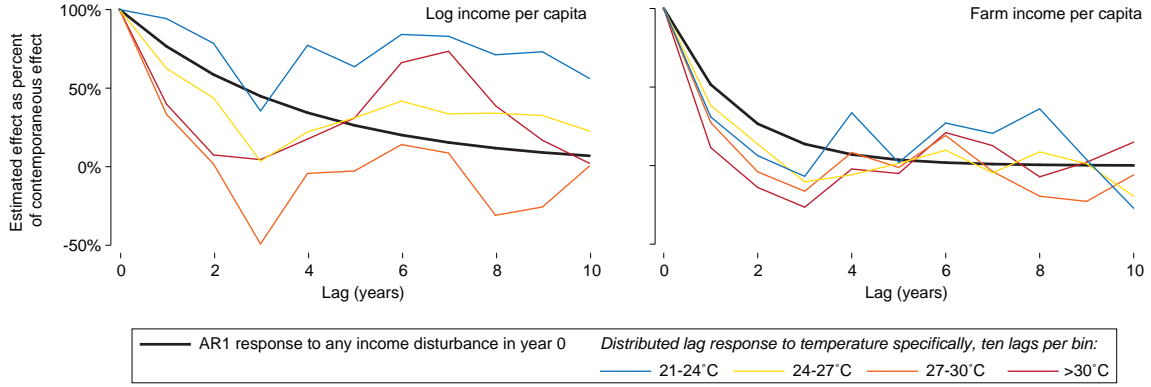


Figure 5: The inter-temporal dynamics of temperature effects on personal income and farm income. All effects are relative to the estimated contemporaneous effect within each temperature bin. Black lines are the effect of any disturbance to county income in the AR1 benchmark model. Colored lines are estimated lagged effects of temperatures in the top four bins, using a model that omits the AR1 structure.

hot day. If the economic disturbance persists indefinitely, that might indicate a permanent change in allocations \mathbf{b} , which could in turn indicate a violation of our assumption that economies are initially at an optimum $\mathbf{b}^*(\mathbf{C})$ determined by their long-run climate. Thus, we examine the persistence of output effects in the years following a marginal change in the temperature distribution.

We also directly examine whether the effects of temperature changes and any endogenous responses to them are more or less persistent than those of other idiosyncratic income changes.²⁴ Estimation of Equation 17 recovers $\hat{\rho} = 0.825$ for log personal income per capita. Thus, an income loss of \$1 in county i in year t will result in a conditional income loss of \$0.825 in year $t + 1$, \$0.68 in year $t + 2$ and so on, relative to the case of no loss in year t . However, this pattern of auto-correlation is largely identified off of variations in the idiosyncratic disturbance term ϵ —not from changes in county temperature.

To isolate the effects of *temperature-induced* changes in income, we estimate a variant of Equation 17 where we replace the lagged dependent variable term ($\rho Y_{i,t-1}$) with ten annual lags of each temperature and precipitation bin. If current temperatures affect future income similarly to other types of income disturbances, then the structure of these lags should be similar to the negative-exponential structure of persistence indicated by $\hat{\rho} Y_{i,t-1}$. The results for log income per capita, as well as for farm income—which we analyze further in Section 6—are shown in Figure 5.²⁵ For comparability, all lags are normalized to the contemporaneous effect of that temperature bin. The thick black line corresponds to the

²⁴We thank James Stock for this suggestion.

²⁵See the Appendix for an analogous test of earnings results.

repeated exponentiation of the estimated autoregressive coefficient ($\hat{\rho}^t$, normalized). Thin colored lines are the lagged effect of a single additional day in the indicated bin.

The effects of hot days clearly decline over time with an average structure that resembles the negative exponential decay captured in our auto-regressive model, although the 252 additional parameters needed to characterize persistence (9 additional lags for 16 temperature bins and 12 precipitation bins) cause the unrestricted lagged coefficients on temperature bins to be much noisier than the baseline model. For total income, the effects of the hottest three bins decay more rapidly than in the baseline model, a response partly captured in the baseline model as the delayed positive response from hot temperatures (recall Figure 4). For farm income, the initial persistence of temperature-driven income changes tends to be substantially lower than in the baseline model, but the ratios for lags larger than three years are very similar to the autoregressive term (and close to zero).²⁶ Nonetheless, in both cases we observe effects that are persistent but do not appear to be permanent, supporting our key assumption that these perturbations are marginal.

Our finding that changes in the current temperature distribution of a county change the future income trajectory of that county roughly similarly to the effects of other contemporaneous income disturbances has two additional implications worth noting. First, this result provides county-level support for prior country-level findings that temporary climatic conditions alter future economic production, an effect modeled using growth rates in earlier work (Hsiang, 2010; Dell, Jones, and Olken, 2012; Burke, Hsiang, and Miguel, 2015).²⁷ Second, evidence of any persistence implies that the total NPV cost of a change in the temperature distribution is larger than the contemporaneous marginal product would suggest. Using a 3% discount rate and $\hat{\rho}$, a change in total income due a shift in the temperature distribution produces an NPV of income changes that is 5 times as large as the contemporaneous effects in Figure 4.²⁸ Below, we do not present this NPV and instead focus on the contemporaneous marginal product of climate, although we do account for persistence in constructing economic projections under climate change.

Stationarity of the relationship over time We consider whether the marginal product of climate has changed appreciably over the time period in our sample (1969-2011). Some consequences of climatic conditions have surely changed over this period, as have many other

²⁶It is possible that the relatively faster recovery of farm income is due to the particular ways in which the PPF in agriculture recovers from temperature changes, although some of this recovery is likely due to crop insurance indemnities paid out as a result of hot temperatures (see the Appendix).

²⁷Prior studies argued that such persistent effects could materialize, for example, if the evolution of durable non-climatic state variables, such as the rate of investment in the capital stock, is influenced by the climate, thereby causing historical climate conditions to affect future periods (Burke, Hsiang, and Miguel, 2015).

²⁸ The NPV of the altered income trajectory is a linear scaling of coefficients by $\frac{1}{1-\hat{\rho}\delta}$: $\sum_{s=t}^{\infty} \delta^{(s-t)} \Delta Y_{is}^m \approx \sum_{s=t}^{\infty} \delta^{(s-t)} \hat{\rho}^{(s-t)} \hat{\beta}^m = \frac{1}{1-\hat{\rho}\delta} \hat{\beta}^m$. Using a discount factor $\delta = 0.97$ (implying an annual discount rate of 3%) and $\hat{\rho} = 0.825$, we estimate this scaling factor to be 5.01.

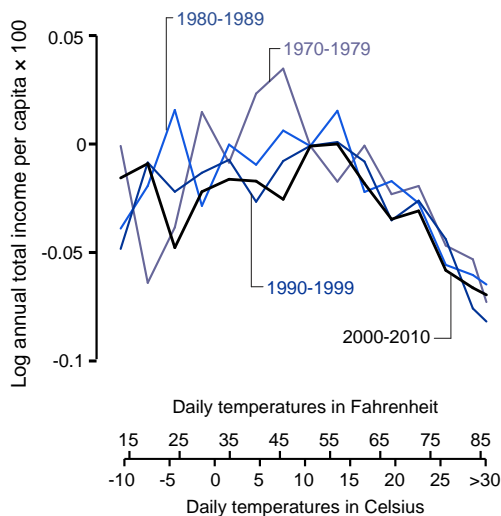


Figure 6: The stationary marginal product of climate over four decades. Figure shows the gradient of the value function for temperature (as in Figure 4) using decade-long subsamples.

aspects of the US economy. For example, the spread of residential air conditioning appears to be responsible for reducing heat-related mortality during the last half-century (Barreca et al., 2016). Meanwhile, other technology, such as the heat-tolerance of maize, remained essentially stationary during this time (Roberts and Schlenker, 2011). Importantly, prior analyses that document how sensitivity to climate evolves over time do not account for the net cost of those adaptations, i.e., opportunity costs due to re-optimization of \mathbf{b}^* , and they cannot account for the (potentially offsetting) substitutions or complementarities that inform the full economic value of these changes in equilibrium. By contrast, our estimate of the gradient in the value function captures all market costs and benefits of adaptations as well as general equilibrium effects, the net effects of which we can observe over time.

To that end, we re-estimate the response of income to temperature separately for each decade. These estimates, shown in Figure 6, are noisier because each relies on a smaller sample, but they do not differ meaningfully from our pooled estimate or from one another. In particular, the marginal effect of warm and hot days is essentially constant over time, demonstrating a remarkable stability of this component of the gradient of the value function. Importantly, such stability does *not* imply that welfare-increasing adaptations to temperature, such as the spread of air conditioning, have not occurred during the last half-century in the US (recall that observable adaptive adjustments $\frac{db}{dC}$ may be arbitrarily large in Equation 12), but it does indicate that the collective net impact of all such adaptations did not fundamentally alter the marginal product of climate. High temperature days remain costly and moderate temperatures remain relatively productive even into the twenty-first century.

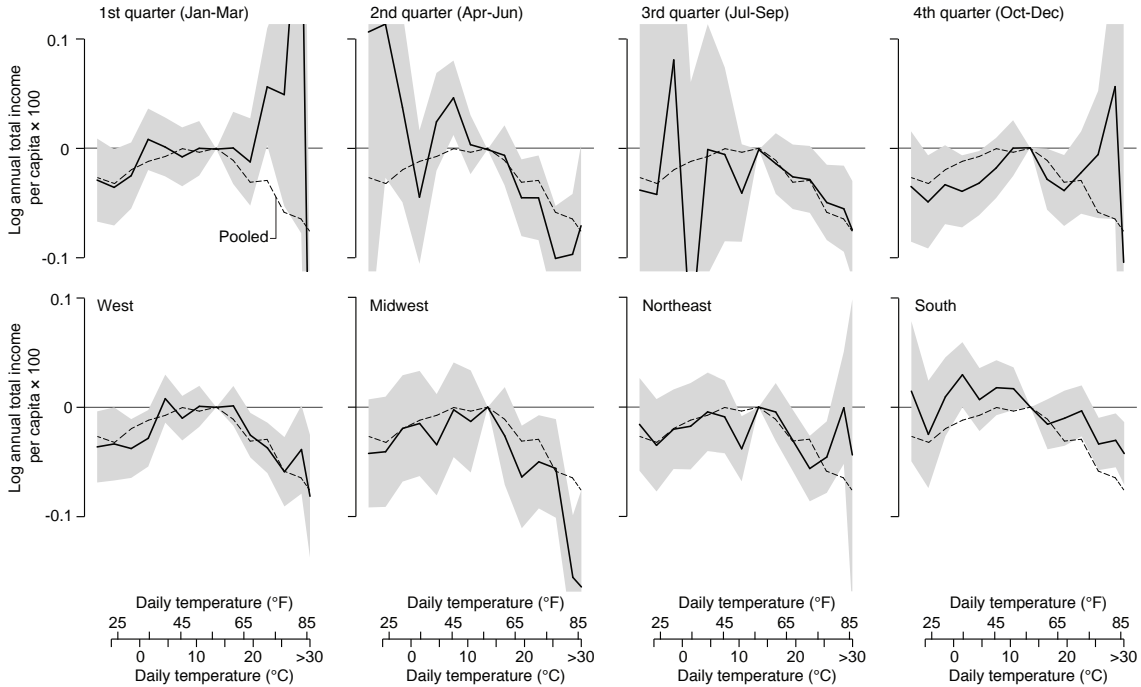


Figure 7: The marginal effect of climate estimated for sub-periods within each year and sub-regions of the country, compared to the pooled estimate (dashed, identical to Figure 4). Top row: Effects of temperatures experienced in each quarter of the year (estimated simultaneously), where annual income is allowed to respond to changes in the temperature distribution for each quarter separately. Bottom row: Effects of annual temperature distributions by region (estimated simultaneously).

Seasonal and regional heterogeneity Next, we examine whether the marginal product of climate varies substantially across seasons by estimating a version of Equation 17 where we construct a set of temperature bins for each quarter. The results are shown in the top row of Figure 7. The marginal effects of warm and hot days in the second and third quarters (roughly corresponding to spring and summer) are extremely similar to each other and to the pooled estimates. A similar pattern holds for cold days in the first and fourth quarters. However, we lack precision to make meaningful comparisons between cold days in the second and third quarters or between hot days in the first and fourth quarters, because such days are uncommon. Overall, these results show that, for those portions of \mathcal{C} with sufficient data, the marginal product of climate is quite consistent across seasons.

Finally, we examine regional heterogeneity in the data. Recall that the marginal product of climate is the *local* gradient of the value function $V(\mathbf{C})$, measured in the vicinity of some position $\mathbf{C}_0 \in \mathcal{C}$. If the local shape of the value function depends on the climate of a location, the local gradient vector will change as \mathcal{C} is traversed, producing *curvature* in the value function. Such curvature could result from endogenous adaptations to different climates

when populations select different $\mathbf{b}^*(\mathbf{C})$ (e.g. investing in air conditioners), as illustrated in Figure 2. Curvature in $V(\mathbf{C})$ could also result from features of the income-generating process that are beyond agents’ control, since the PPF may itself respond nonlinearly to temperature (e.g. Schlenker and Roberts, 2009) as might aggregate demand (e.g. Auffhammer, Baylis, and Hausman, 2017). For any of these reasons, if there is such curvature, then the marginal product of climate will be different for different regions of the country, because counties in different regions will, on average, be located in different parts of \mathcal{C} .

We examine the marginal product of climate for four major regions of the country, the West, Midwest, South, and Northeast (as defined by the US Census), searching for *prima facie* evidence of curvature in $V(\mathbf{C})$. The results, shown in the bottom row of Figure 7, indicate that the overall structure of the response to daily temperature distributions is similar across the country, suggesting that curvature in the value function may exist but is not dramatic for most temperatures. No regional subsample exhibits a response that is statistically different from the pooled estimate at any temperature, although the structure of point estimates at high temperatures provides suggestive evidence that the effect of such temperatures is not identical everywhere. In particular, high temperatures are most costly in the Midwest and least costly in the South. In Section 6 we investigate possible reasons for these differences, but first we consider the structure of curvature in the value function directly.

5.2 Results for a value function with curvature

We now present results that do not assume the gradient vector of the value function $V(\mathbf{C})$ is constant, effectively allowing the marginal effect of distorting the temperature distribution by \mathbf{C}' to depend on the county’s initial position within \mathcal{C} . Accounting for curvature in the value function becomes important if, for example, the marginal effect of additional hot days declines as a population experiences and adapts to higher numbers of hot days. To estimate a curved surface, we set $H = 3$ in Equation 17 such that income is cubic in the count of days in each bin, allowing the value function to curve in any of its sixteen dimensions. Because this approach fully accounts for changes in the marginal effect of temperature distributions due to adaptation, we denote it the “full adaptation model,” in contrast to the previously estimated “constant marginal effect” model.

Figure 8 presents the results of these two models for all temperature bins above 15°C. Because the gradient vector of the value function in the full adaptation model is no longer constant, it cannot be plotted as in earlier sections. Instead, we display the *total* (non-marginal) contribution of days within each temperature bin to annual income. Each curve can be thought of as a one-dimensional cross-section through the 16-dimensional value function, with the derivative representing the *marginal* effect of an additional day at that temperature. For each temperature bin, each individual county will be roughly centered

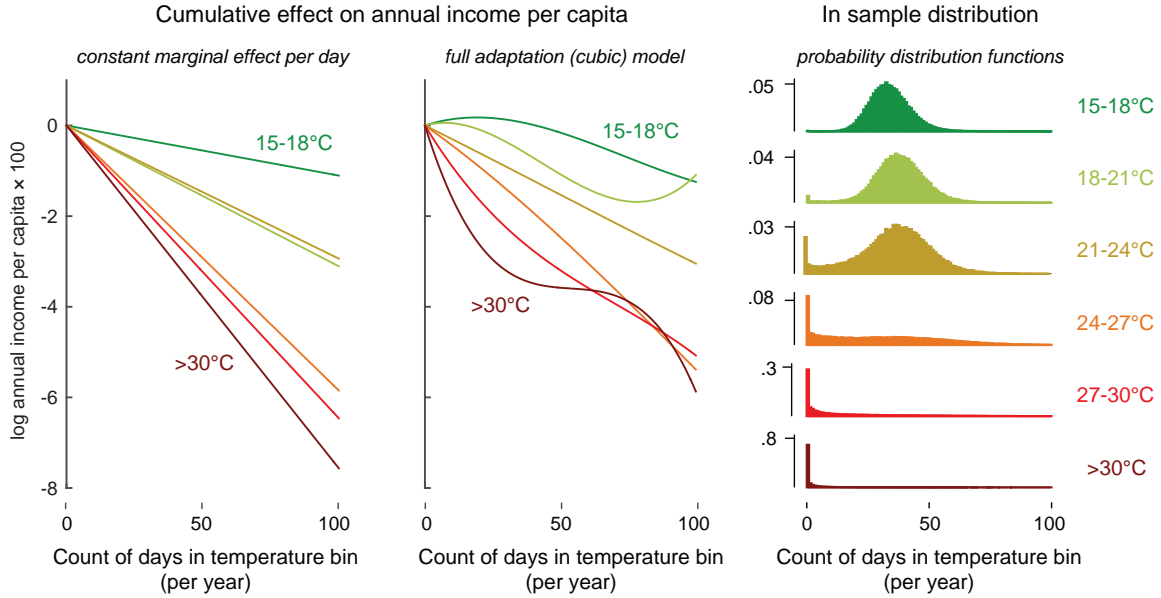


Figure 8: Comparison of the estimated effect of hot temperature days using an approach that assumes constant marginal effects of days in each temperature bin ($H = 1$, left) and an approach that allows income to be a cubic function of the count of days in each temperature bin ($H = 3$, middle). Histograms (right) display the empirical distribution of counts in each temperature bin across all county-years in our sample. Temperature bin colors are fixed across all three panels.

at its average count of days in that bin (its climate), with small variations in that count (weather) contributing to our estimate of the local slope of the corresponding curve at that position. For reference, the right-most panel displays empirical probability distributions over the count of days in each temperature bin across our sample of county-years, using the same color scheme as the center and left panels.

In the constant marginal effects model (left-most panel of Figure 8), annual income declines linearly in the count of days within each warm or hot temperature bin, with slopes that are equal to the constant marginal effects in Figure 4. In the full adaptation model (middle panel of Figure 8) the marginal effects of warm temperatures (15-18°C and 18-21°C) are less negative for small counts of days, becoming more negative as counts of days increase, while the marginal effect of very hot temperatures (27-30°C and > 30°C) are most negative for small counts of days. For the “middle-hot” temperature bins (21-24°C and 24-27°C) the marginal effects of additional days appear to be essentially constant.

For the 27-30°C bin, the convexity is modest, with the first day in this bin reducing annual incomes by 0.096%, the tenth day reducing income by 0.081%, and the thirtieth day reducing income by 0.054%. Because almost all county-years experience fewer than

ten days in this category, the in-sample response remains well-approximated by a linear function. Only for the very hottest bin, where temperatures exceed 30°C , does curvature in the value function have such a large effect that the affine approximation is poor. Having one such hot day lowers annual income by 0.181% (recall that the average day contributes roughly 0.274% to annual income); the tenth hot day a year lowers annual income by only 0.125% and the thirtieth lowers income by a still smaller 0.039%. It is worth noting that over 70% of county-years in our sample have only one or zero days in this bin (right panel of Figure 8), and only very hot regions of the country have a large number of days with such high temperatures.²⁹ The flattening out of the value function indicates that these areas make large adaptive allocations in \mathbf{b}^* such that hot days have limited marginal impact.

To our knowledge, this estimate of the non-affine 16-dimensional value function represents the first characterization of the market value of temperature, accounting for all benefits and costs of adaptive adjustments captured by the market. Importantly, our results reveal that, for most regions in \mathcal{C} currently populated by modern US counties, the constant marginal effects model is a strikingly good approximation of the curved surface described by the complete adaptation model. This insight, which we could not have assumed *ex ante*, is powerful because it allows us to more confidently exploit the constant marginal effects model when exploring additional heterogeneity in the data, simplifying our analysis and description of the value function. However, the adaptive responses captured by curvature in the value function become increasingly relevant as populations move into regions of \mathcal{C} that contain large numbers of hot and very hot days, such as in our projection of future climate changes impacts. Thus, if marginal effects are assumed to be constant in these simulations, we will mis-estimate the total effect of warming the climate, which we demonstrate in Section 7.2.

Note on the sign of forecast bias A “folk theorem” salient in the climate-economics literature states that climate change projections using econometric estimates that assume constant marginal effects of temperature represent an “upper bound” for the damage from warming, since un-modeled adaptations will cause the actual marginal damages from warming to be smaller. Figure 8 demonstrates that this idea originates from correct theoretical intuition, in the sense that populations with many hot days adapt to them, but the reasoning is incomplete when applied to regression analyses. Focusing on the response to very hot ($> 30^{\circ}\text{C}$) days for simplicity, we see that the constant marginal effects model *does* over-estimate the marginal damage from additional warming relative to the full adaptation model for highly adapted populations with many (e.g. 40) very hot days—this is consistent with the intuition of the “folk theorem.” However, the constant marginal effects model also

²⁹There are extremely few county-year observations with more than 50 days above 30°C , causing the slope of the surface above this cutoff to not be statistically different from zero.

under-estimates the marginal damage of warming for cooler and poorly adapted populations that presently have few very hot days (e.g. 1). The slope of the fully adapted response is actually *steeper* at one day above 30°C (middle panel) than the constant marginal effects estimate (left panel). Because the constant marginal effects model recovers the pooled average treatment effect in the sample, many cooler locations—which in the US represent the vast majority of counties—are assigned marginal effects that are too small in magnitude. Thus, as these numerous counties with very few hot days warm up, they initially descend down the fully adapted curve much more rapidly than they would descend down the constant marginal effects line. As we demonstrate in Section 7.2, because most economic activity in the present-day US occurs in these relatively cool counties, fully accounting for adaptation causes the projected losses from warming to *increase* relative to projections that use a constant marginal effects model. The error in the “folk theorem” originates from failing to account for the influence of *historical* adaptive behaviors, which are conceptually identical to the future adaptations it assumes will occur but are already present in modern data.

6 Mechanisms

Next, we try to explore some of the mechanisms that underlie the results presented above using two strategies. First, we explore the response of different sectors, since sectoral PPFs may respond to temperature. Second, we stratify counties based on the extent to which they have made prior investments in durable assets that are thought to be substitutes for temperature in the economy, namely air-conditioning and urban infrastructure.

6.1 Farm, non-farm, and manufacturing income

In the US, high daily temperatures are known to reduce yields of major crops (Schlenker and Roberts, 2009)³⁰ as well as to reduce labor productivity among farm workers (Stevens, 2017) and labor supply among workers exposed to outdoor temperatures (Graff Zivin and Neidell, 2014), which includes manufacturing workers. These responses will necessarily alter a county’s PPF and thus could, in principle, explain our main finding. However, these studies alone are not conclusive. If local or regional prices change with temperature, perhaps in response to concurrent changes in aggregate demand, then changes in production quantities might not translate into changes in revenue. We cannot observe local prices or aggregate demand directly, but we can investigate whether temperature-dependence of the PPF may be contributing to the overall effect of temperature on income by separately examining its agricultural and non-agricultural components.

To examine how crop losses contribute to our main result, we repeat our analysis with *log (revenue from crop sales per capita)* as the dependent variable. The response function,

³⁰Also see Mendelsohn, Nordhaus, and Shaw (1994); Schlenker, Hanemann, and Fisher (2005); Deschênes and Greenstone (2007); Fisher et al. (2012); Burke and Emerick (2016).

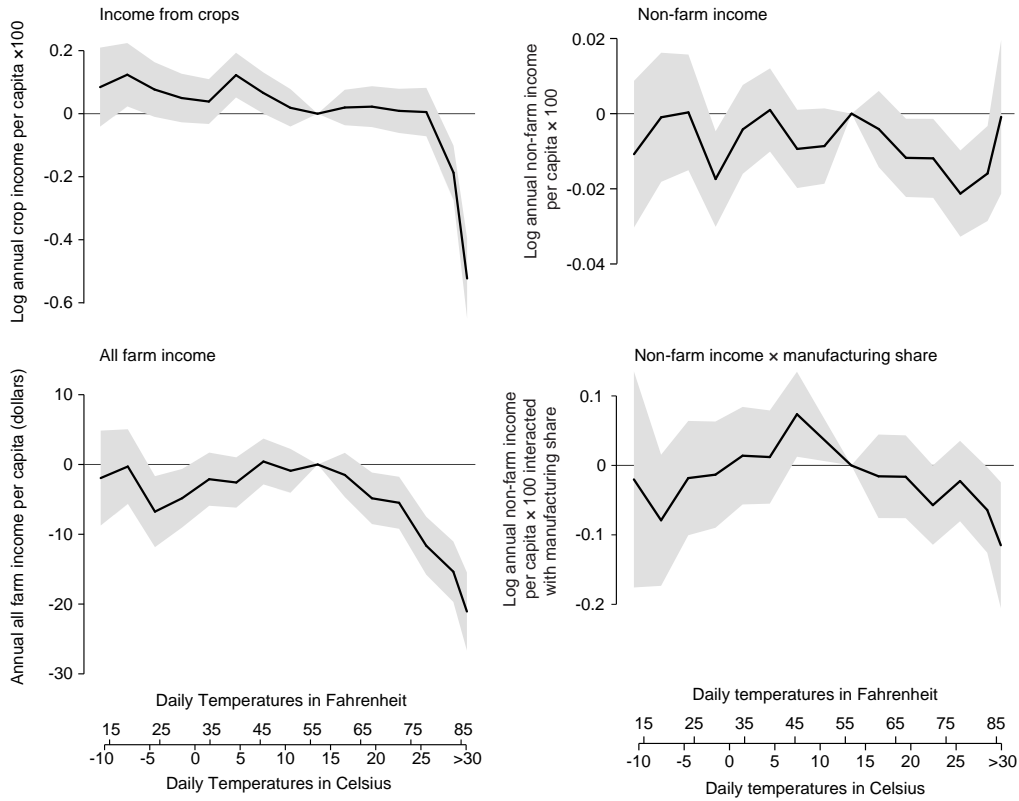


Figure 9: The marginal effect of daily average temperature on log income from crops per capita, net farm income per capita (in levels), log non-farm income per capita, and log non-farm income per capita interacted with county-level measures of manufacturing income share (U.S. Census Bureau, 1969-2011).

shown in the top left of Figure 9, exhibits steep declines when daily average temperatures rise above 27°C. This structure is quantitatively similar to the yield response estimated by Schlenker and Roberts (2009) (see Appendix Figure A6 and the accompanying discussion for a detailed comparison). Our results suggest that higher crop prices do not meaningfully offset yield losses caused by high temperatures, perhaps because regional crop markets are well-integrated and local yield losses do not substantively alter prices. Thus, reductions in county crop yields as estimated by Schlenker and Roberts (2009) translate essentially one-for-one into revenue reductions.

In the bottom left panel of Figure 9, we examine how *net farm income per capita* in levels³¹ responds to daily temperature and estimate that it declines \$21.07 for each day above 30°C. The structure of the response differs somewhat from that of crop income: we observe a fall in net farm income starting at temperatures around 20°C. We lack the data

³¹Net farm income is not analyzed in logs because many observations are negative.

to test mechanisms that could explain these differences at lower temperatures. However, the observed structure is broadly consistent with the labor productivity response discussed below and in Stevens (2017). Alternatively, farmers may be increasing expenditure on inputs to combat the negative impacts of moderate temperatures on yields. Finally, it is also possible that other productive factors, such as livestock, respond negatively to these temperatures.

We examine the role of temperature in the generation of non-farm income by using *log non-farm income per capita* as the dependent variable (top right of Figure 9). Non-farm income is relatively flat (albeit noisy) at low temperatures and then declines systematically at temperatures above 15°C, the same breakpoint observed for total income. However, the percentage effect on non-farm income is smaller in magnitude. For example, temperatures at 25°C lower annual non-farm income by only 0.021% relative to 15°C, whereas the analogous loss of annual total income is 0.058%. These non-farm income estimates are broadly consistent with the labor supply response documented by Graff Zivin and Neidell (2014) (see Appendix for a detailed comparison) as well as with non-agricultural labor productivity responses from factories and lab experiments (e.g. Parsons, 2014).

Finally, in part motivated by a century of experiments that implicate temperature in influencing manufacturing worker productivity (Huntington, 1922; Parsons, 2014), we estimate how the marginal effect of temperature on non-farm income evolves with changes in a county’s manufacturing share. Specifically, we use County Business Patterns (U.S. Census Bureau, 1969-2011) to calculate the ratio of manufacturing payroll to total payroll (in 2011 dollars) over the time period 1969–2011. We then add an interaction term between each temperature bin and manufacturing payroll share to our baseline specification to identify a component of temperature-related non-farm income variation that projects systematically onto the spatial distribution of manufacturing. The interacted effects, shown in the bottom right of Figure 9, indicate an inverted-U shaped relationship between manufacturing income and daily temperature that peaks around 9-12°C (48.2-53.6°F) and declines roughly 1% per 1°C of warming, the estimated effect of temperature if all non-farm payroll were in manufacturing. These results indicate that the net negative effects of cold or hot temperatures on non-farm income tend to increase with a county’s average manufacturing share, perhaps due to previously documented effects on the manufacturing PPF via labor productivity. Notably, the lower optimum could also in part be driven by the effects of temperature on non-labor factors or on factor reallocations within or across industries.

6.2 Evidence of specific substitutes for climate in production

A complementary strategy for understanding the determinants of the marginal product of temperature is to examine how allocation decisions embodied by \mathbf{b} influence the observed marginal product. If populations adjust \mathbf{b}^* in response to their temperature distribution

while maximizing production, then some factors in \mathbf{b} must be effective substitutes for certain dimensions of \mathbf{C} . Optimizing populations will substitute human-made capital for “natural capital” (Hartwick, 1977; Solow, 1991), where natural capital in this context are dimensions of \mathbf{C}_i .

To that end, we look for empirical evidence of such behavior by examining how the marginal product of temperature changes based on the allocation of specific factors thought to be substitutes for climate.³² Specifically, we consider the effectiveness of two potential substitutes to climate in the production process: air conditioning and urbanization. The potential of the former to substitute for climate is intuitive. Urbanization is a more complex phenomenon that is surely not driven purely or even primarily by the desire to adapt to climate. Nonetheless, urbanization is likely to alter the effects of climate on the production process by altering the organization, density, and composition of economic activity in such a manner that it becomes less affected by temperature (Kahn, 2013; Deschênes et al., 2011).

Air conditioning We classify the counties in our sample into three groups based on residential air conditioning (AC) penetration rates reported in the 1980 Census:³³ (1) 60% or less, (2) 60–80%, and (3) 80% or more. The last group consists almost exclusively of counties in Florida, Kansas, Louisiana, Oklahoma, and Texas. We then estimate a version of Equation 17 where we interact indicators for each of these three groups with contemporaneous temperature and precipitation bins.

The results for total income per capita, farm income per capita, and non-farm income per capita are shown in Figure 10. Counties with AC penetration rates below 60% and 60%–80% show similar susceptibility to high temperatures with respect to total income, both to each other and to the pooled sample (Panels A–B). Counties with the highest AC penetration rates, on the other hand, appear to be half as susceptible to such temperatures, although these estimates are noisier (Panel C).

The next six panels show separate results for farm and non-farm income. Farm income responses vary somewhat across these groupings (Panel D–E), but all decline at hot temperatures regardless of AC penetration, presumably because AC does not directly benefit agriculture. By contrast, non-farm income in counties with the highest AC penetration exhibit essentially no response temperature (Panel I), while counties with lower AC penetration experience declines in non-farm income with higher temperatures (Panels G–H).

Urbanization Next, we classify counties in our sample as “urban” if a majority of their population lived in an urban area in 2010, as reported by the U.S. Census Bureau.³⁴ 41%

³²This exercise examines the gradient of the value function when specific elements in \mathbf{b} are held fixed.

³³This is last complete cross-sectional survey of AC ownership in the US. See Barreca et al. (2016) for a detailed discussion.

³⁴To our knowledge, this statistic is not available in earlier years.

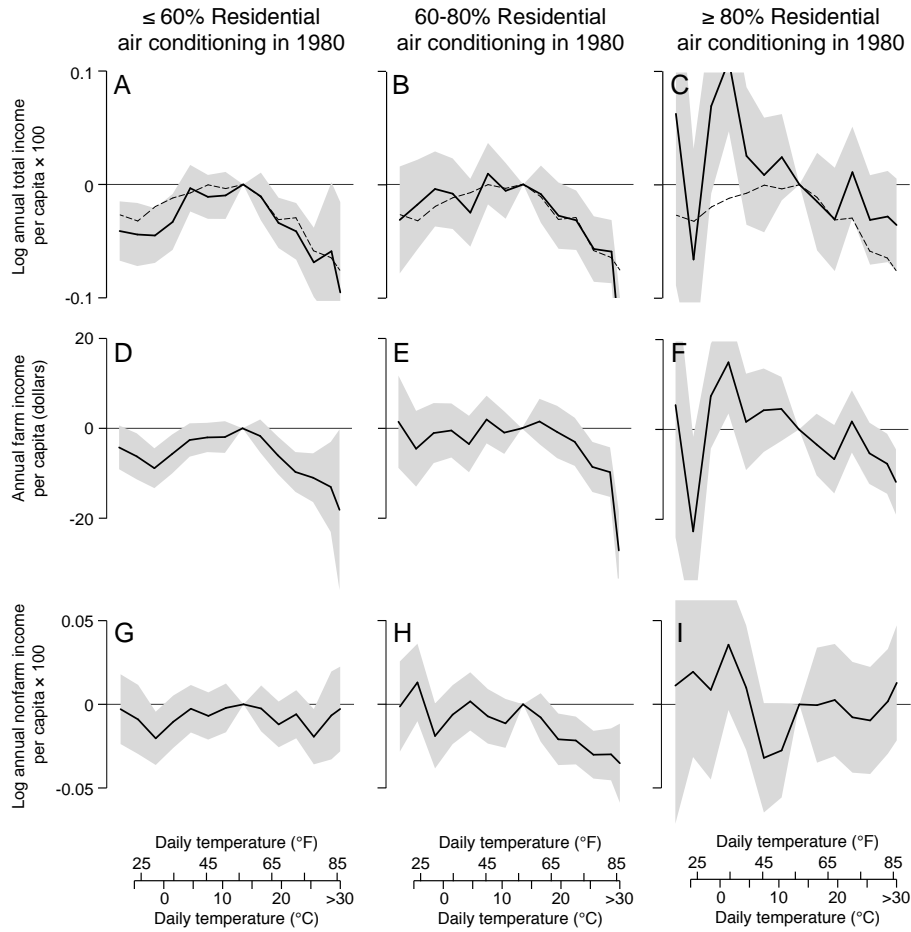


Figure 10: Estimated effects on various income measures by 1980 residential air conditioning penetration, as reported by the 1980 Census. Dashed lines correspond to the pooled estimate from Figure 4. Results for all subsamples (within a single row) are jointly estimated. Shaded areas are 95% confidence intervals.

of U.S. counties are “urban” by this definition. As with AC penetration, we interact the contemporaneous temperature and precipitation bins with urban/rural indicators to estimate two value functions: one that applies to urban counties and one that applies to all other counties. The results for log total income per capita are shown in Figure 11. Rural counties are slightly more susceptible than average to the hottest temperatures, although the differences are not statistically significant (left panel). By contrast, point estimates for urban areas suggest a lower susceptibility to heat (center panel).

Additionally, we experiment with an “ultra-urban” category for counties with population density above the 90th percentile (averaged over the sample). Counties at the top of the population density distribution also exhibit income losses at warm and hot temperature (right panel), although effects at the highest temperatures are somewhat reduced.




Figure 11: Effect of temperature on income by urbanization status. Results for left and center panels are jointly estimated. Results for right panel are jointly estimated with counties below the 90th percentile of population density (not shown). Shaded areas are 95% confidence intervals. Dashed lines correspond to the pooled estimate from Figure 4.

Overall, our results suggest that AC penetration is a more important predictor of income sensitivity to heat than urbanization, although both matter. Because allocation of resources toward AC is primarily a direct response to warmer temperatures, the costs and benefits of such allocations will be implicitly captured by the “full adaptation model” that allows the value function to curve. In the full adaptation model, the reduction of income-sensitivity with AC adoption is captured by the reduction in marginal effects for hot locations (recall Figure 8), to the extent that these investments are driven by climatic conditions. It is less obvious that the influence of urbanization is explicitly captured in the previously estimated “full adaptation” value function. Rather, our main estimates simply reflect average treatment effects integrated over the current (uncorrelated) spatial distribution of urbanization and climates. Thus, to account for the effects of all endogenous adaptations to temperature—including AC—as well as for the influence of urbanization in our valuations of the current and future climate (below), we estimate a “full adaptation” version of the value function that contains separate curved surfaces for both urban and non-urban counties.

7 Valuing current and future climates

The stability of our estimates for expected income suggests that most, if not all, county-level economies lie on a coherent $V(\mathbf{C})$ surface offset by a county-specific constant μ_i , which captures all other county-specific factors. By applying our estimates for the marginal product of climate $\frac{dV(\mathbf{C})}{d\mathbf{C}}$ along a line integral through the climate space \mathcal{C} (Equation 15), we can trace out a non-marginal change in income for county i if the climate were displaced from an initial value \mathbf{C}_{i1} to an alternative climate \mathbf{C}_{i2} , net of all adaptive adjustments in \mathbf{b}_i^*

(which we do not need to observe directly) along the path:

$$\Delta Y_i = \int_{\mathbf{C}_{i1}}^{\mathbf{C}_{i2}} \frac{d\widehat{V}(\mathbf{C})}{d\mathbf{C}} d\mathbf{C},$$

where either \mathbf{C}_{i1} or \mathbf{C}_{i2} is the current climate and the other is the counterfactual.

We compute ΔY for two counterfactual climates. First, we consider how much current daily temperature distributions contribute to current production by displacing the climate of a single benchmark county ($\mathbf{C}_{i1} = \mathbf{C}_0$), which provides a common starting point, to the actual observed climate of each county ($\mathbf{C}_{i2} = \mathbf{C}_i$). Second, we gradually distort the current climate of each county along a “business as usual” climate change scenario and compute how its productivity changes (relative to its historic climate) through 2100. Throughout, we use estimates from the autoregressive model. As a robustness check, we have replicated our valuation of the future climate replacing the autoregressive term by 9 additional weather lags, as in Figure 5. These loss estimates, available upon request, are generally larger.

7.1 Contribution of temperature to current production

To understand how current temperatures contribute to cross-sectional productivity patterns, we integrate our fully non-linear estimate of $\frac{d\widehat{V}(\mathbf{C})}{d\mathbf{C}}$ from the benchmark historical climate of Lebanon, Kansas (\mathbf{C}_0) to the historical climate of each US county (\mathbf{C}_i), both averaged over 1968-1990. We choose Lebanon simply because it is the geographic centroid of the country, making the results of this comparison easy to visualize; the *relative* contributions of temperature climates to counties’ economies do not depend on the benchmark county.³⁵

The left panel in Figure 12 depicts the difference in income between each county and Lebanon, KS (marked with a black circle) that is attributable to differences in daily temperature distributions. Colors represent the income change a county would exhibit as the climate of Lebanon is smoothly transitioned to the climate of each county in our pooled “full adaptation” model. Thus, the map depicts the “height” of the value function evaluated at historical climates across the country, after all county-specific differences (μ_i) and adaptive adjustments (\mathbf{b}^*) have been accounted for. Note that while the underlying model uses only *marginal within-county* variation in daily temperature distributions, it yields clear and sensible *non-marginal between-county* differences regarding the value of their climates. The white band of counties, stretching through the corn belt, south of Appalachia, and up to the Mid-Atlantic states, indicates temperature climates that are similar in economic value to Lebanon. South of this band, the value of climate declines as the number of low-productivity hot days increases, with locations along the Gulf Coast losing more than \$1000 per capita (2011 dollars) in annual income, relative to Lebanon. Note that these values fully account

³⁵The choice of benchmark county is unimportant because it is a constant of integration that is differenced out in cross-county comparisons of ΔY_i .

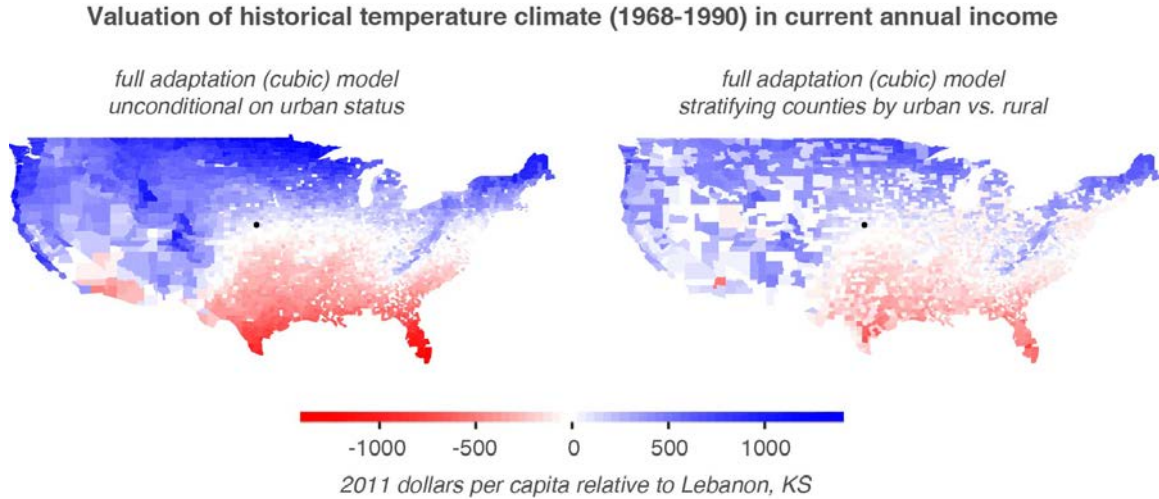


Figure 12: Value of historical county climates relative to the climate of Lebanon, Kansas (black circle). The left map uses the third-degree polynomial pooled model. The right map uses a third-degree polynomial model interacted with a county’s urban status.

for the curvature in $V(\mathbf{C})$, which includes changes in the marginal effect of additional hot days due to adaptation. North of the white zero-band, counties earn higher incomes due to their climate, largely due to the reduction or elimination of hot days, with the climate of locations along the Pacific, Rockies, Great Lakes, and New England generating \$500-1000 or more per capita annually.³⁶

The right panel of of Figure 12 shows an analogous calculation, but one that conditions the value of the climate on whether a county is urban or not. Using this approach, regional patterns are largely unchanged, but, consistent with Figure 11, urban locations exhibit a more muted version of regional patterns. The distribution of urban counties does not appear systematically correlated with climatic conditions, so we find it useful to think of this installed urban capital as being determined by some orthogonal optimization that we do not observe. The pooled model (left panel) can thus be considered the *a priori* valuation of each county’s climate if the urban-rural status of each county were unknown, whereas the stratified model (right panel) is a valuation where the status of current urban investments is known in advance.

7.2 Production distortions due to future warming

Next, we use 44 different climate change scenarios from Hsiang et al. (2017), constructed to emulate the probability distribution of global climate sensitivities, to project how output will change due to future warming in RCP8.5 (“business as usual”) relative to a coun-

³⁶As discussed earlier in footnote 28, accounting for the dynamic effects of these income losses would increase their magnitude by a factor of 5.

terfactual where temperatures remain at historical levels. For each county, the scenarios report the expected number of days in each 1-degree-Celsius temperature bin in 2080-2099. We aggregate this distribution to the 3-degree temperature bins used in our estimates and use the empirical 1969-1990 distribution of temperatures in each county as the no-climate-change counterfactual. We assume that warming begins in 1991, is linear in the number of days in each temperature bin, and converges to the 2080-2099 distribution in 2090 since it is the midpoint of the interval. This approach smoothly and realistically transitions daily temperature distributions while maintaining a total count of 365 days in each year of the projection. We use three discount rates (1%, 3%, and 5%) to probe the sensitivity of the projections to this important parameter and calculate the net present value (NPV) of lost income relative to the no-climate-change scenario. We multiply the per-capita estimates by the county’s actual ($t \leq 2011$) or projected population ($t > 2011$), assuming county-specific linear population growth.

We apply the warming projections to the constant marginal effects model and the full adaptation model, in both cases pooling all counties together. These two simulations allow us to see the effect of accounting for adaptation costs and benefits. To account for the estimated substitutability of climate and current urbanization, which is not correlated with the temperature and thus not accounted for in the full adaptation model, we also create projections using the full adaptation model interacted with an urban indicator and current patterns of urbanization are held fixed into the future.

The spatial distribution of NPV for the median climate trajectory (in terms of the total income loss) is shown in the first column of panels in Figure 13. The units are billions of 2011 dollars. Rows indicate the specification used to estimate the value function. Here, we only display the 3% discount rate case. Spatial patterns are essentially unchanged with different discount rates, although magnitudes become more exaggerated or muted. Without accounting for non-linearity or heterogeneity (linear model), the largest aggregate losses from climate change appear concentrated in the Southwest and the Northeast (dark red). Some Northern states and many counties in Florida also suffer large losses, and very few counties see income gains (light and dark blue). However, once we allow the value function to be curved (full adaptation model), more counties, especially in Texas, are projected to experience income gains as a result of climate change; allowing for urban-rural heterogeneity produces yet more gains in urban counties in Gulf Coast states. Consistent with the earlier discussion of the erroneous “folk theorem,” allowing for curvature in the value function increases damage projections for initially cool counties, such as the Northeast and Midwest. This occurs because the value function for these counties becomes *steeper* when the relatively flat marginal effects of Southern states are no longer pooled with these cooler locations (recall Figure 8). Because economic production is more heavily concentrated

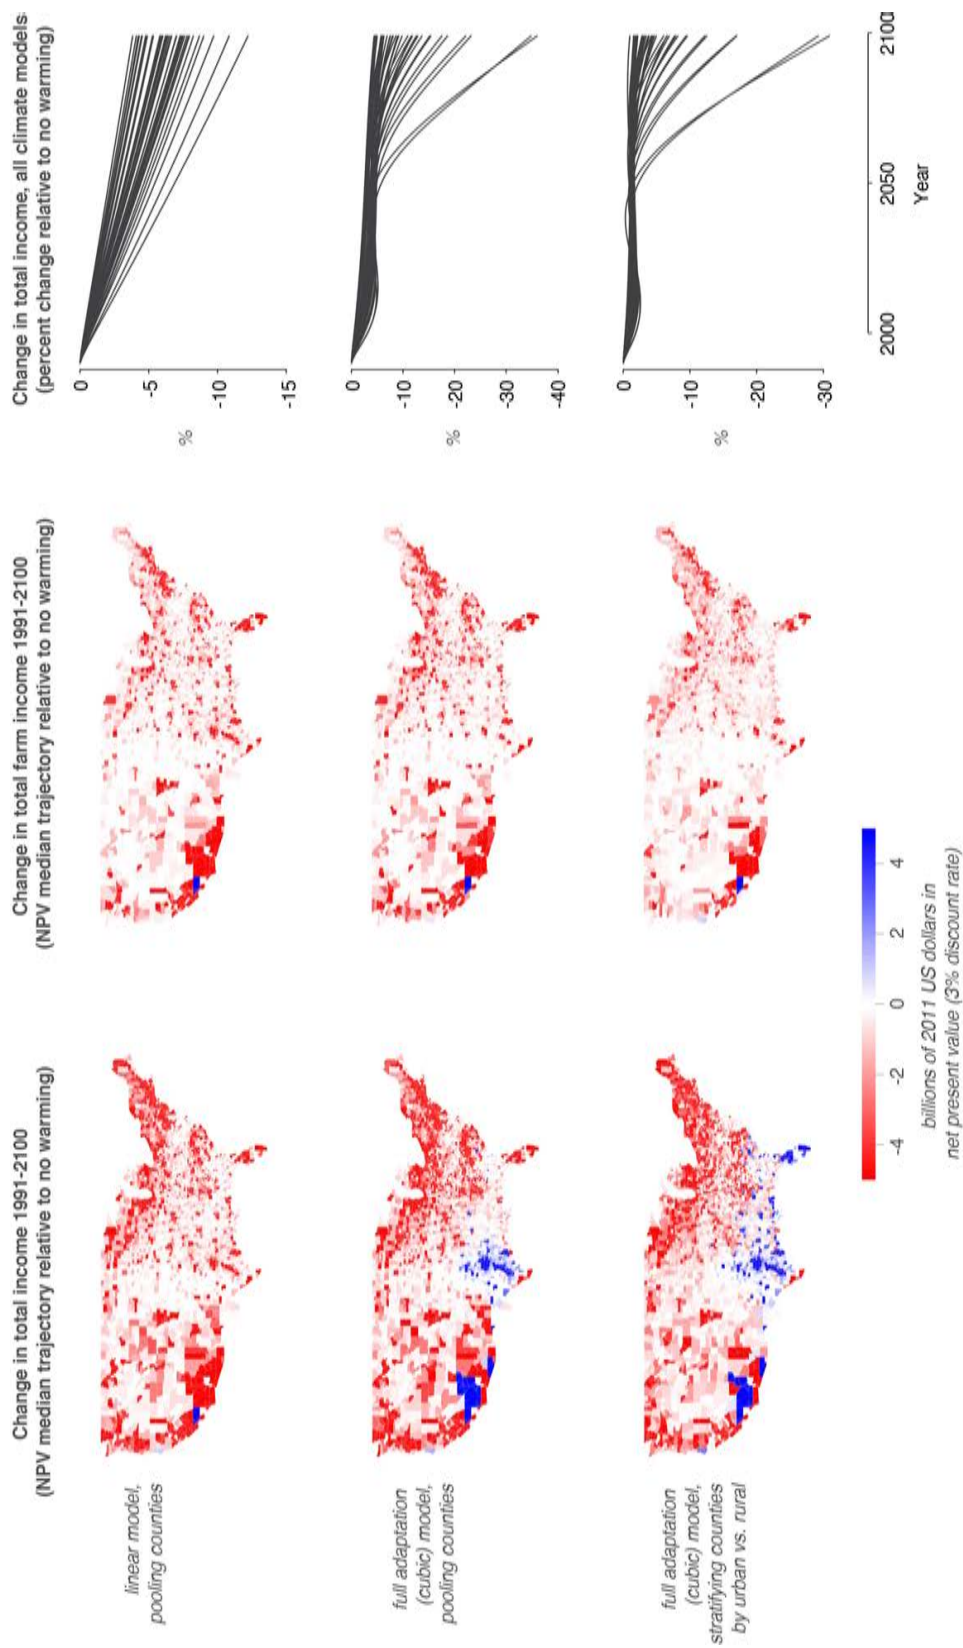


Figure 13: Outcome variable specified above each column. Projections in the rightmost column combine estimates from models specified by row headings with 44 model projections for 2080-2099 from Hsiang et al. (2017). Maps show the spatial distribution of the net present value of income losses for the median climate projection, as measured by the estimated impact in 2099.

Table 1: **Probability distribution of the net present value of national income losses due to “business as usual” climate change (RCP8.5) through the year 2099**

	(1)	(2)	(3)
Discount rate:	1%	3%	5%
Panel A: constant marginal effects model, no heterogeneity			
10th percentile	-43.54	-18.93	-13.55
25th percentile	-37.44	-16.33	-11.7
Median	-31.52	-13.7	-9.8
75th percentile	-26.5	-11.5	-8.22
90th percentile	-22.19	-9.61	-6.86
Panel B: “full adaptation” cubic model, no heterogeneity			
10th percentile	-60.41	-25.37	-18.17
25th percentile	-46.44	-20.94	-15.43
Median	-38.96	-17.87	-13.3
75th percentile	-31.48	-15.55	-11.61
90th percentile	-26.01	-11.98	-8.82
Panel C: “full adaptation” cubic model, urban-rural heterogeneity			
10th percentile	-27.32	-10.36	-7.15
25th percentile	-18.26	-7.9	-5.79
Median	-13.47	-6.74	-5.16
75th percentile	-10.62	-5.35	-4.1
90th percentile	-9.11	-4.65	-3.51

Discount rate shown above each column. Values are in trillions of US 2011 dollars.

in the North, allowing for adaptation by letting marginal damages vary by climate has the net effect of *increasing* total national income losses.

In contrast, projected changes in farm income (second column of panels in Figure 13) are extremely similar across the three sets of estimates—consistent with our finding and the findings of others that some adaptation technologies, such as AC, do not play a major role in agriculture (e.g. Roberts and Schlenker, 2011; Burke and Emerick, 2016). Very few counties are projected to gain agricultural income as a result of climate change. The largest losses are again concentrated in the Southwest, Northeast, and Florida. Comparing the two

columns, we see that while farm income losses contribute to total income losses, they do not fully explain their magnitude, especially once heterogeneity and non-linearity are accounted for. Non-agricultural losses thus play a major role in the projected reductions in economic output, consistent with prior country-level results (Burke, Hsiang, and Miguel, 2015).

The third column of Figure 13 shows how aggregate income losses (as a percent of annual income) evolve over time across the 44 climate projections. By definition, income losses are zero in 1990. In the linear model, losses grow linearly between 1991 and 2099 and range from 3.8% to 12.2% of the no-climate-change counterfactual income in 2099. The median scenario predicts income losses of 6.3% in 2099. Both projections based on cubic estimates (with and without urban-rural heterogeneity) display non-linear and non-monotonic patterns and a wider distribution of projected income losses by 2099. Without accounting for urban-rural heterogeneity, 2099 losses range from 4.4% to 36% of income (median is 8.0%). Accounting for such heterogeneity reduces the magnitude of the lower and upper bound of losses to 0.98% and 31%, respectively, and shifts the median to 3.4%.

In Table 1 we present a summary of the distribution of the NPV of aggregate income losses across all 44 climate projections using each of the three models for the value function and three different annual discount rates (1%, 3%, and 5%). Units are in trillions of US 2011 dollars. Panel A shows estimates using the constant marginal effects model. With a 3% discount rate, the median NPV of income loss is \$13.7 trillion (2011 dollars). At the 10th percentile of the distribution, losses are almost \$19 trillion, and at the 90th percentile losses are estimated at \$9.6 trillion (higher percentiles are more positive). The NPV of losses is more than two times larger when we use a discount rate of 1% rather than 3%. Conversely, using a rate of 5% yields losses that are about one-quarter to one-third lower.

Accounting for adaptation by allowing for a cubic relationship between income and the number of days in a temperature bin (Panel B) increases income loss estimates. As discussed above, this occurs because the marginal effects of warming for the high temperature bins are negatively correlated with counties' overall economic output. However, adding urban-rural heterogeneity while still allowing for a cubic relationship (Panel C) reduces projected losses, yielding the smallest loss estimates. This adjustment substantially reduces total costs because economic activity is concentrated in urban counties, and allowing for heterogeneity reduces marginal damage from warming in these counties. Specifically, the median NPV of income losses is \$6.7 trillion at a 3% discount rate and ranges from \$4.7 trillion at the 90th percentile to \$10 trillion at the 10th percentile. Varying the discount rate affects the estimates in Panels B and C similarly to Panel A.

8 Discussion

Previous analyses of how an economy is influenced by its climate have struggled to simultaneously account for unobservable factors that differ across locations and the overall impact, net of costs, of adaptive adjustments. Here we developed a general reformulation of the problem that delivers both with respect to the role of the climate in economic production. Applying this approach to estimate the role of daily temperature distributions for market output in the modern US, we recover the marginal product of climate and demonstrate how it can be integrated to recover the overall impact of non-marginal climate changes.

There are several limitations to our analysis. First, our results depend on competitive markets being efficient in the long run, following Arrow and Debreu (1954). For example, we assume agents have perfect information about the climate they inhabit, that capital can be rented at annualized costs, and that if there are profitable opportunities they will be seized. Without these assumptions, our Envelope Theorem result, which depends on firms optimizing and markets clearing, may no longer hold exactly. Future work should examine how market distortions, imperfect information, and the incomplete rationality of decision-makers may alter these findings.

It is also important to note that our estimates capture only the impact of shifting *temperature* distributions. Any present or future influence of other climatic factors—such as storm frequencies, sea levels, or drought—are omitted. In principle, it is straightforward to extend this approach to additional dimensions of the climate, something we view as an avenue for future work.

We also assume that disturbances due to weather are “small” such that they do not move the economy “too far” from its equilibrium. This assumption guarantees that the Envelope Theorem holds, and its validity depends on the spatial and temporal scale of analysis as well as how weather and climate relate to the economy. We chose to demonstrate our approach using annual distributions of daily temperature, described with temperature bins, in part because perturbing an annual temperature distributions by shifting a few days from one bin to the next is plausibly a “small” perturbation in an otherwise large space of possible temperature distributions.³⁷ This notion is confirmed by our finding that shifting individual days results in only fractional and temporary changes in annual percentages of income. Yet this assumption might not necessarily hold for all dimensions of climate—for example, hurricanes and mega-droughts may not be sufficiently “small” economic perturbations for our approach to be applicable (Hornbeck, 2012; Deryugina, Kawano, and Levitt, forthcoming).

Importantly, our results do not represent welfare effects. Our focus is characterizing the

³⁷The size of this space, in terms of dimension, can be controlled by using more or fewer bins.

contribution of the temperature climate to aggregate economic production, as captured by total market revenue. Other analyses consider the economic value of the climate in welfare terms (e.g. by using an Integrated Assessment Model) and/or account for non-market effects, such as increasing economic inequality, degraded ecosystems, higher crime, or the loss of life (Anthoff, Hepburn, and Tol, 2009; Hsiang et al., 2017). These factors may be substantially affected by the temperature climate but are not accounted for in the present analysis beyond any influence they have on the structure of the PPF or aggregate demand.

One of the contributions of our analysis is to account for all of the allocative adjustments made *within* a macroeconomy to cope with a change in the same county’s climate. Our empirical implementation focuses on a large number of “small” macroeconomies, US counties, within which a large number of possible allocations exist. There is, however, interest in allocative adjustments *across larger spatial scales*, such as the reallocation of production or labor (via migration) across regions (Desmet and Rossi-Hansberg, 2015; Costinot, Donaldson, and Smith, 2016; Dingel, Hsiang, and Meng, 2017; Desmet et al., 2017). Our theoretical analysis is also directly relevant for these larger scales, as the domain of the macroeconomy under consideration can be expanded to contain a larger region with no other adjustment needed (e.g. Burke, Hsiang, and Miguel, 2015).

Finally, it is crucial to acknowledge that future, unknowable technological innovations that may affect the marginal product of climate are unlikely to be captured in our empirical analysis. Theoretically, new technologies can be incorporated into our framework by increasing the dimensionality of \mathcal{B} to allow for allocations towards a new type of technology (which has infinite cost prior to discovery). However, it is not possible for us to empirically explore the structure of the value function in the subspace of \mathcal{B} corresponding to a not-yet-existent technology. Nonetheless, future innovations might be captured to the extent that they are represented in the present marketplace. For example, valuations of assets (e.g. stocks, land values) reflect market beliefs about the future trajectory of technology and affect allocation decisions within the present market. Additional work should explore the extent to which current allocations may be informative about the path of future technologies and their potential role in altering the marginal product of climate.

References

- Anthoff, David, Cameron Hepburn, and Richard SJ Tol. 2009. “Equity weighting and the marginal damage costs of climate change.” *Ecological Economics* 68 (3):836–849.
- Aroonruengsawat, Anin and Maximilian Auffhammer. 2011. “Impacts of climate change on residential electricity consumption: evidence from billing data.” In *The Economics of climate change: Adaptations past and present*. University of Chicago Press, 311–342.
- Arrow, Kenneth J and Gerard Debreu. 1954. “Existence of an equilibrium for a competitive economy.” *Econometrica* :265–290.
- Auffhammer, Maximilian, Patrick Baylis, and Catherine H Hausman. 2017. “Climate change

- is projected to have severe impacts on the frequency and intensity of peak electricity demand across the United States.” *Proceedings of the National Academy of Sciences* .
- 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 Deschênes, Michael Greenstone, and Joseph S. Shapiro. 2016. “Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century.” *Journal of Political Economy* .
- Burke, Marshall, John Dykema, David B Lobell, Edward Miguel, and Shanker Satyanath. 2015. “Incorporating Climate Uncertainty into Estimates of Climate Change Impacts.” *Review of Economics and Statistics* .
- Burke, Marshall and Kyle Emerick. 2016. “Adaptation to Climate Change: Evidence from US Agriculture.” *American Economic Journal: Economic Policy* 8 (3):106–40.
- Burke, Marshall, Solomon Hsiang, and Edward Miguel. 2015. “Global non-linear effect of temperature on economic production.” *Nature* 527 (7577):235–239.
- Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller. 2011. “Robust inference with multiway clustering.” *Journal of Business & Economic Statistics* 29 (2).
- Carleton, Tamma A and Solomon M Hsiang. 2016. “Social and economic impacts of climate.” *Science* 353 (6304):aad9837.
- Costinot, Arnaud, Dave Donaldson, and Cory Smith. 2016. “Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world.” *Journal of Political Economy* 124 (1):205–248.
- 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.
- . 2013. “What Do We Learn from the Weather? The New Climate-Economy Literature.” Tech. rep., National Bureau of Economic Research.
- Deryugina, T. 2017. “The Fiscal Cost of Hurricanes: Disaster Aid Versus Social Insurance.” *American Economic Journal: Economic Policy* .
- Deryugina, Tatyana, Laura Kawano, and Steven Levitt. forthcoming. “The Economic Impact of Hurricane Katrina on its Victims: Evidence from Individual Tax Returns.” *American Economic Journal: Applied Economics* .
- Deschênes, Olivier, Robin Burgess, David Donaldson, and Michael Greenstone. 2011. “Climate Change and Mortality in India.” *Mimeo* .
- 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.
- Deschênes, Olivier and Michael Greenstone. 2011. “Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US.” *American Economic Journal: Applied Economics* 3 (4):152–85.
- Desmet, Klaus, Robert Kopp, David Nagy, Michael Oppenheimer, and Estban Rossi-Hansberg. 2017. “Dynamic Geography of the Population and Economic Response to Sea Level Rise.” *Mimeo* .
- Desmet, Klaus and Esteban Rossi-Hansberg. 2015. “On the spatial economic impact of global warming.” *Journal of Urban Economics* 88:16–37.
- Dingel, Jonathan, Solomon Hsiang, and Kyle Meng. 2017. “The Spatial Structure of Endowments, Trade, and Inequality: Evidence from the Global Climate.” *Mimeo* .
- 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.

- 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.
- Guo, Christopher and Christopher Costello. 2013. "The value of adaption: Climate change and timberland management." *J. of Env Econ and Management* 65 (3):452–468.
- 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 Hsiang, Michael Delgado, Amir Jina, Kate Larsen, Michael Mastrandrea, Shashank Mohan, Robert Muir-Wood, DJ Rasmussen, James Rising, and Paul Wilson. 2015. *American Climate Prospectus: Economic Risks in the United States*. Columbia University Press.
- Hsiang, Solomon. 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.
- . 2016. "Climate econometrics." *Annual Rev of Resource Econ* 8:43–75.
- Hsiang, Solomon, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, DJ Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, Kate Larsen, and Trevor Houser. 2017. "Estimating economic damage from climate change in the United States." *Science* 356 (6345):1362–1369.
- Hsiang, Solomon and Daiju Narita. 2012. "Adaptation to Cyclone Risk: Evidence from the Global Cross-Section." *Climate Change Economics* 3 (2).
- Hsiang, Solomon M and Marshall Burke. 2014. "Climate, conflict, and social stability: what does the evidence say?" *Climatic Change* 123 (1):39–55.
- Huntington, Ellsworth. 1922. *Civilization and climate*. Yale University Press.
- Kahn, M. 2013. *Climatopolis: how our cities will thrive in the hotter future*. Basic Books.
- Kelly, David L, Charles D Kolstad, and Glenn T Mitchell. 2005. "Adjustment costs from environmental change." *Journal of Environmental Econ and Management* 50 (3):468–495.
- Koopmans, Tjalling C. 1957. "Allocation of resources and the price system." In *Three Essays on The State of Economic Science*, chap. 1. Oxford: McGraw-Hill, 4–40.
- 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 Rev* :753–771.
- Nickell, Stephen. 1981. "Biases in dynamic models with fixed effects." *Econometrica* :1417–1426.
- Parsons, Ken. 2014. *Human thermal environments: the effects of hot, moderate, and cold environments on human health, comfort, and performance*. CRC press.
- Roberts, Michael J. and Wolfram Schlenker. 2011. *The Evolution of Heat Tolerance of Corn: Implications for Climate Change*. University of Chicago Press, 225–251.
- Schlenker, W, W Hanemann, and A Fisher. 2005. "Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach." *American Economic Review* 95 (1):395–406.
- Schlenker, Wolfram, W Michael Hanemann, and Anthony C Fisher. 2006. "The impact of global warming on US agriculture: an econometric analysis of optimal growing conditions." *Review of Economics and statistics* 88 (1):113–125.
- 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. 1991. “Sustainability: an economist’s perspective.” In *Economics of the Environment*, edited by Robert Stavins. W. W. Norton & Company.
- Stevens, Andrew. 2017. “Temperature, Wages, and Agricultural Labor Productivity.” .
- U.S. Census Bureau. 1969-2011. “County Business Patterns.” Available from <https://www.census.gov/programs-surveys/cbp/data.html>.

Appendix (For Online Publication)

This appendix supplements the main text with the following six sections:

- **A1 - Relationship of this formulation to prior work**
- **A2 - Example value function in a two-dimensional climate**
- **A3 - Data Appendix**
- **A4 - Additional features of the economic response to temperatures**
- **A5 - Quantitative comparison of agricultural and non-agricultural effects on the PPF**
- **A6 - Point estimates for key results**

A1 Relationship of this formulation to prior work

Seminal analysis by Mendelsohn, Nordhaus, and Shaw (1994) attempted to directly estimate Equation 8 in a cross-sectional nonlinear regression of farm profits on a vector $\hat{\mathbf{C}}$ that captured average seasonal temperatures and rainfall. This approach essentially specifies that average temperatures and rainfall are sufficient statistics to reconstruct, through application of $\psi(\cdot)$, the full distribution of actual weather \mathbf{c} relevant to farm value. If, conditional on observable characteristics included in the regression (such as soil quality), farms are identical and only the first moments of temperature and rainfall are relevant to output, then this approach will recover the shape of $V(\mathbf{C})$ net of all adaptation costs and benefits.

Schlenker, Hanemann, and Fisher (2006) expanded on this cross-sectional approach by adopting a more sophisticated structure for $\psi(\cdot)$, whereby degree-days above and below two specified temperature cutoffs are considered sufficient statistics $\hat{\mathbf{C}}$ for estimation of $V(\mathbf{C})$.

Deschênes and Greenstone (2007) raise the concern that that different farm units may not be comparable, even conditional on observable traits, leading to potential bias in these earlier regression frameworks. To circumvent this issue, they propose to use a within-unit panel regression approach that differences out any constant unobserved heterogeneity between farm units. To implement this, the authors assume $\hat{\mathbf{C}} = \mathbf{c}_{it}$ and then estimate a version of Equation 8 exploiting random variation in \mathbf{c}_{it} . In their implementation, the authors used first moments in temperature and rainfall to summarize weather, analogous to Mendelsohn, Nordhaus, and Shaw (1994). A concern raised by later authors was that endogenous responses to climate changes captured by re-optimization of $\mathbf{b}^*(\mathbf{C})$ would not be captured in this framework, since farmers can differentiate between temporary changes in \mathbf{c}_{it} and long-term changes in \mathbf{C} .

Schlenker and Roberts (2009) and Deschênes and Greenstone (2011) expanded on the approach of Deschênes and Greenstone (2007) by adopting sophisticated structures for $\psi(\cdot)$ to capture nonlinear responses of crop yields and human mortality, respectively, to temperature. These contributions did not directly address adjustments of $\mathbf{b}^*(\mathbf{C})$.

Analyses by Aroonruengsawat and Auffhammer (2011), Hsiang and Narita (2012), and Barreca et al. (2016), along with others, built on these contributions by accounting for some re-optimization of $\mathbf{b}^*(\mathbf{C})$ in a panel framework where the partial effect of \mathbf{c}_{it} on the outcome, in an analog to Equation 8, is estimated directly, allowing this effect to vary as a function of \mathbf{C} —thereby capturing some influence of $\mathbf{b}^*(\mathbf{C})$ by proxy. While this approach is able to document the presence of adaptive behaviors, it has now been recognized that it cannot fully capture changes in $V(\mathbf{C})$ because the costs of adjusting factors \mathbf{b} is unobserved by the econometrician (Houser et al., 2015).

Dell, Jones, and Olken (2012) and Burke and Emerick (2016) also expand on the approach of Deschênes and Greenstone (2007) by using a long (multi-year) period of observation t when constructing $\hat{\mathbf{C}} = \mathbf{c}_{it}$, arguing that the period is sufficiently long that $\mathbf{b}^*(\mathbf{C})$ would have plausibly adjusted. Neither analysis recovers evidence of such adjustment, concluding that such adjustments are absent. However, even if evidence of adjustment had been found, it would not be possible to evaluate the costs (and thus net benefits) of these adaptations.

Thus, a systematic challenge to evaluating the economic value of climate has been the inability to simultaneously account for unobservable heterogeneity while also accounting for adaptive re-optimization of \mathbf{b}^* in a manner that fully accounts for both costs and benefits (Hsiang, 2016). We solve this challenge in a single framework by carefully constructing the appropriate $\psi(\cdot)$, allowing for nonlinear adaptation at all points in the distribution of $f_{\mathbf{x}}(\cdot)$, and restricting our analysis to an optimized outcome where short-run marginal changes in \mathbf{c}_{it} exactly identify the marginal effect of long-run changes in \mathbf{C} .

A2 Example value function in a two-dimensional climate

We extend the graphical depiction of the value function along a North-South road with a one-dimensional climate ($K = 1$, developed in Section 3 and shown in Figure 2) to a two-dimensional climate ($K = 2$). We continue to have only one dimension of adaptation ($N = 1$). Expansion of the climate space by one dimension is useful for developing intuition for how this approach generalizes as the dimensionality of \mathcal{C} increases. Our actual empirical implementation for the US explores a 16-dimensional climate with unknown N , making it more difficult to visualize.

Consider a $\psi(\cdot)$ such that the probability distribution of daily temperatures in a year is summarized by a three-bin histogram, shown in Panel A of Figure A1. C_1 is the expected

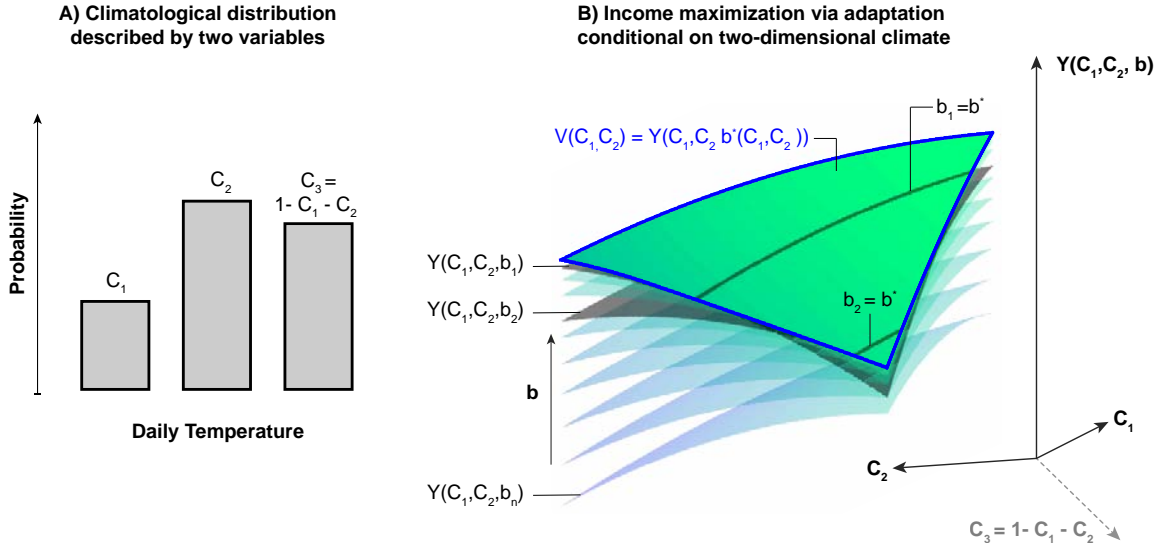


Figure A1: Application of the Envelope Theorem to valuing climate. (A) An example climate defined by a three temperature bin histogram, where the probability distribution of outcomes is fully defined by the mass in two bins (C_1 and C_2). (B) Translucent surfaces are production surfaces over the space (C_1, C_2) for different values of the control variable b . Production maximizes $b = b^*(C_1, C_2)$, such that the opaque triangle (outlined in blue) is the maximum output for each climate position after adaptation, defining the value function $V(C_1, C_2)$. Black surfaces highlight regions of V where two individual surfaces represent the maximized quantities and are exactly tangent to V (i.e. via the Envelope Theorem).

fraction of days with temperature below a cutoff \bar{T} . C_2 is the expected fraction of days with temperature above \bar{T} and below a second cutoff temperature $\bar{\bar{T}}$. C_3 is the expected fraction of days with temperature above $\bar{\bar{T}}$ and is fully determined by the first two dimensions of the climate since $C_3 = 1 - C_1 - C_2$. The climate vector is therefore $\mathbf{C} = (C_1, C_2)$. The space of possible climates is then the unit 2-simplex:

$$\mathcal{C} = \{(C_1, C_2) \mid C_1, C_2 \in [0, 1], C_1 + C_2 \in [0, 1]\}.$$

Let there remain only one dimension of possible adaptive adjustment $\mathbf{b} = b$.

Depicting $Y(\mathbf{C}, \mathbf{b}) = Y(C_1, C_2, b)$ now requires four dimensions. Panel B of Figure A1 depicts multiple semi-translucent surfaces, each a function over the 2-simplex \mathcal{C} , holding a value of b fixed. The height of the n^{th} surface at a point (C_1, C_2) is the level of output the economy would exhibit for the climate (C_1, C_2) if $b = b_n$. Optimization of output would lead to selection of $b = b^*(C_1, C_2)$ for each position in \mathcal{C} , causing the actual economy to exhibit production that matched the highest surface at each position, corresponding with the opaque blue-green curved triangle surface that is outlined in blue. This two-dimensional

Table A1: Summary statistics for key variables

	(1)	(2)	(3)	(4)	(5)
	Mean	Std. Dev.	Min	Max	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.

surface is the value function

$$V(C_1, C_2) = Y(C_1, C_2, b^*(C_1, C_2))$$

and it is the two-dimensional analog to the one-dimensional blue ridge-line in Figure 2 of the main text. This curved triangle is the upper envelope of all the production frontiers across all values of b , where individual surfaces of $Y(C_1, C_2, b_n)$ are exactly tangent to the value function for those positions in the climate space (C_1, C_2) where $b_n = b^*$, as was described by Equation 13. For example, the black translucent surface labeled $Y(C_1, C_2, b_2)$ lies below V for almost all positions in \mathcal{C} , but defines the maximum value obtainable for the small band labeled $b_2 = b^*$, where the surfaces are exactly tangent. This tangency is the result of Equation 13, which allows the shape of the value-function to be measured locally by exploiting two-dimensional perturbations (ξ_1, ξ_2) caused by weather without consideration for any adaptive adjustment of b .

A3 Data Appendix

Here, we discuss our data in more detail. Table A1 presents summary statistics for the key variables in our sample. Recall that wage and salary data in REIS are largely derived from employers reporting wage and salary payments for the purposes of unemployment insurance. There are only five industries that are not fully subject to such reporting requirements: 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 an unemployment insurance 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.

In the REIS data, 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.

For agricultural income, there exist states where 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 temperature on farm income should be viewed as a lower bound.

We also obtained data on total transfers from government to individuals from the REIS, the analysis of which is relegated to the next section of this Appendix. These transfers include unemployment insurance, which in turn consists primarily of standard state-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, 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 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 administrative data, obtained through a Freedom of Information Act request.

A4 Additional features of the economic response to temperatures

Here we consider the marginal effect of daily temperature distributions on earnings, transfers, and the effect of temperature in neighboring counties.

Earnings Earnings make up the majority of personal income. In Figure A2 we display the effect of daily temperature on earnings per capita in current and prior years. Qualitatively, the structure of the earnings response is very similar to the total income response in Figure 4, although the magnitudes of the point estimates are larger for earnings. Relative to a day at 15°C (59°F), a day at 29°C (84.2°F) lowers annual earnings by roughly 0.11%. Assuming

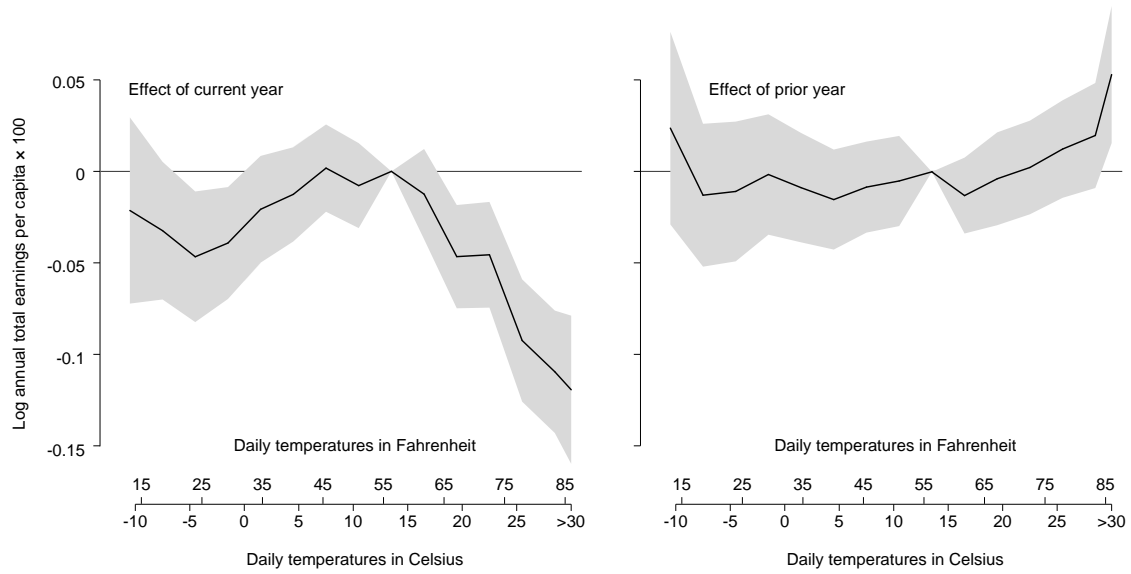


Figure A2: Same as Figure 4 in the main text, except for log total earnings per capita.

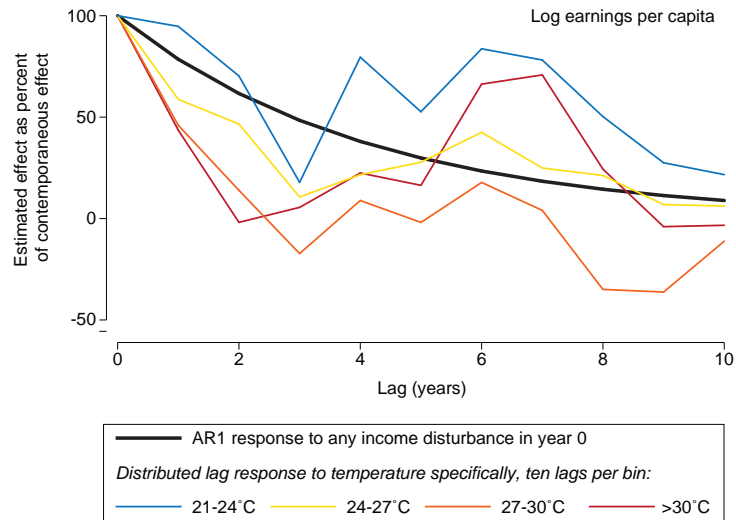


Figure A3: Same as Figure 5 in the main text, except for log total earnings per capita.

uniform output across 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.

Figure A3 shows the structure of temperature lags relative to the AR1 response for log earnings per capita (analogous to Figure 5 in the main text). As before, we focus on the top four temperature bins and normalize the contemporaneous effects to 100. The results are very similar to those for log total income per capita: the temperature effect decays somewhat faster over time relative to other idiosyncratic income shocks, but overall the lag structures look fairly similar.

Transfers from government Prior studies have found that federal government transfers increase following natural disasters (Healy and Malhotra, 2009; Deryugina, 2017), but whether temperature changes lead to a systematic change in the distribution of transfers from the government generally is unknown. In particular, it is plausible that transfers might offset some income losses due to temperature such that we do not observe them—if such masking is occurring it would cause us to under-estimate the marginal product of temperature in our main analysis. To examine whether government transfers might be affecting our estimates, we obtain multiple types of data on transfers, including various types of unemployment insurance, Medicare, federal crop insurance, and *ad hoc* disaster transfers directed by Congress, as described earlier in the Appendix. Figure A4 shows that daily temperatures have no effect on county-level annual transfers from the government (excluding crop-related payments) or on county-level spending on public medical benefits. However, we find evidence that *ad hoc* crop disaster payments increase as a results of very hot days (> 30°C), while crop insurance payouts increase steeply for days that exceed 27°C (80.6°F). The latter estimates suggest that county-level farm income losses would be roughly 25% higher if crop insurance were not available.

Spatial displacement Our model can be generalized to include information on neighbor’s realizations $\mathbf{x}_{i+1,t}$ in the construction of county climate \mathbf{C}_i or weather \mathbf{c}_{it} , if such measures are economically relevant to production Y at i . One plausibly important mechanism through which this could occur is if economic activity is displaced from one county into neighboring counties when temperatures change. Thus, we examine whether there is any evidence for spatial displacement of economic activity across county borders due to temperature by estimating a spatial lag model where income is regressed on the average count of county-days in each temperature bin within five 100km-wide annuli surrounding each county (in addition to all controls and temperature measures in our baseline model). The results for hot temperature bins, shown in Figure A5, indicate that high temperatures

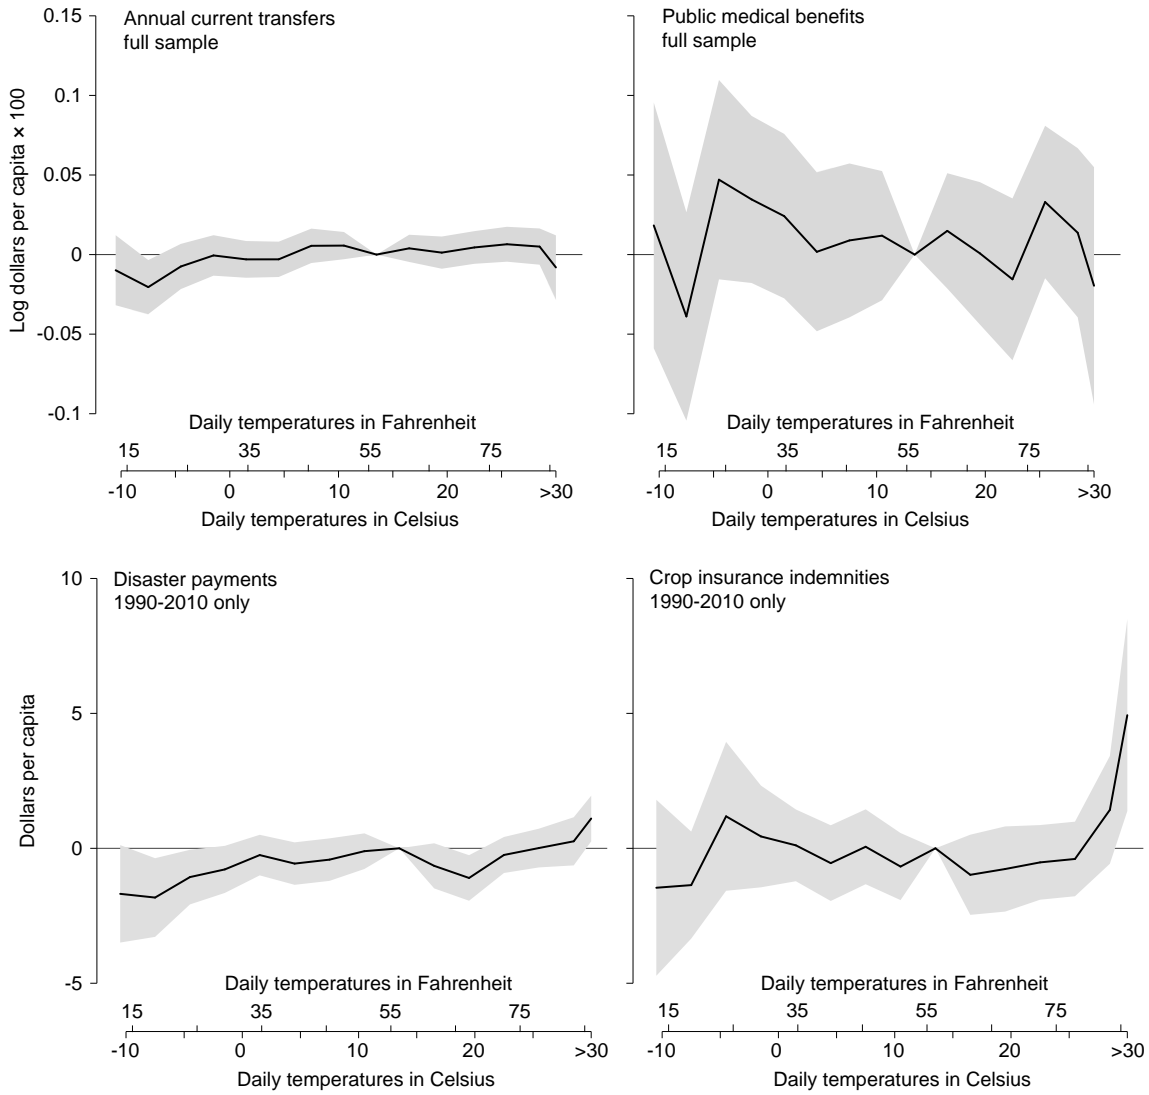


Figure A4: Top left: the effect of daily average temperatures on log total government transfers to individuals per capita. Top right: same, but for the subset of transfers that are public medical benefits. Bottom left: the effect of daily temperatures on *ad hoc* crop disaster payments per capita (in levels). Bottom right: the effect of daily temperatures on crop insurance indemnities per capita.

continue to have a negative effect on own income even when accounting for neighbors' temperature. Further, we do not recover evidence of spatial displacement. If anything, there is some evidence that high temperatures in neighboring counties have negative effects on a county's own income, either because of negative spillovers across counties that experience high temperature days or because 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.

As noted in Section 8 in the main text, one approach for addressing the possibility of any spatial displacement of economic activity is to increase the spatial-scale of aggregation such that any displacement remains contained within the macroeconomy under consideration. The primary drawback of such an approach is that it limits the number of observations available and the extent to which the data span \mathcal{C} . A benefit of such an approach is that at larger scales of aggregation, the Envelope Theorem will hold with increasing accuracy, as allocative decisions within the economy become increasingly continuous.³⁸ Theoretically, there is no reason why the entire world economy could not be utilized as the unit of analysis, thereby capturing the net effect of all possible spatial displacements.

A5 Quantitative comparison of agricultural and non-agricultural effects on the PPF

Here, we compare both the structure and magnitude of farm and non-farm incomes' responses to temperature (see Figure 9) 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 Figure A6 (left and right panel, respectively).

It is worth noting here that Graff Zivin and Neidell (2014) obtain data on the quantity of labor supplied but cannot observe labor effort, i.e. the productivity of labor supplied. Almost a century of 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; Huntington, 1922; Seppanen, Fisk, and Lei, 2006; Parsons, 2014). Thus, the total labor effects on income may be larger than the estimates in Graff Zivin and Neidell (2014) suggest, but the overall structure of the response should be similar.

Crop and Farm Income The decline in crop income explains a significant share (but not all) of our main result for total income: Figure 9 indicated that 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 how temperature effects are measured. The estimated effect in Schlenker and Roberts (2009) is

³⁸We thank Michael Roberts for pointing this out.

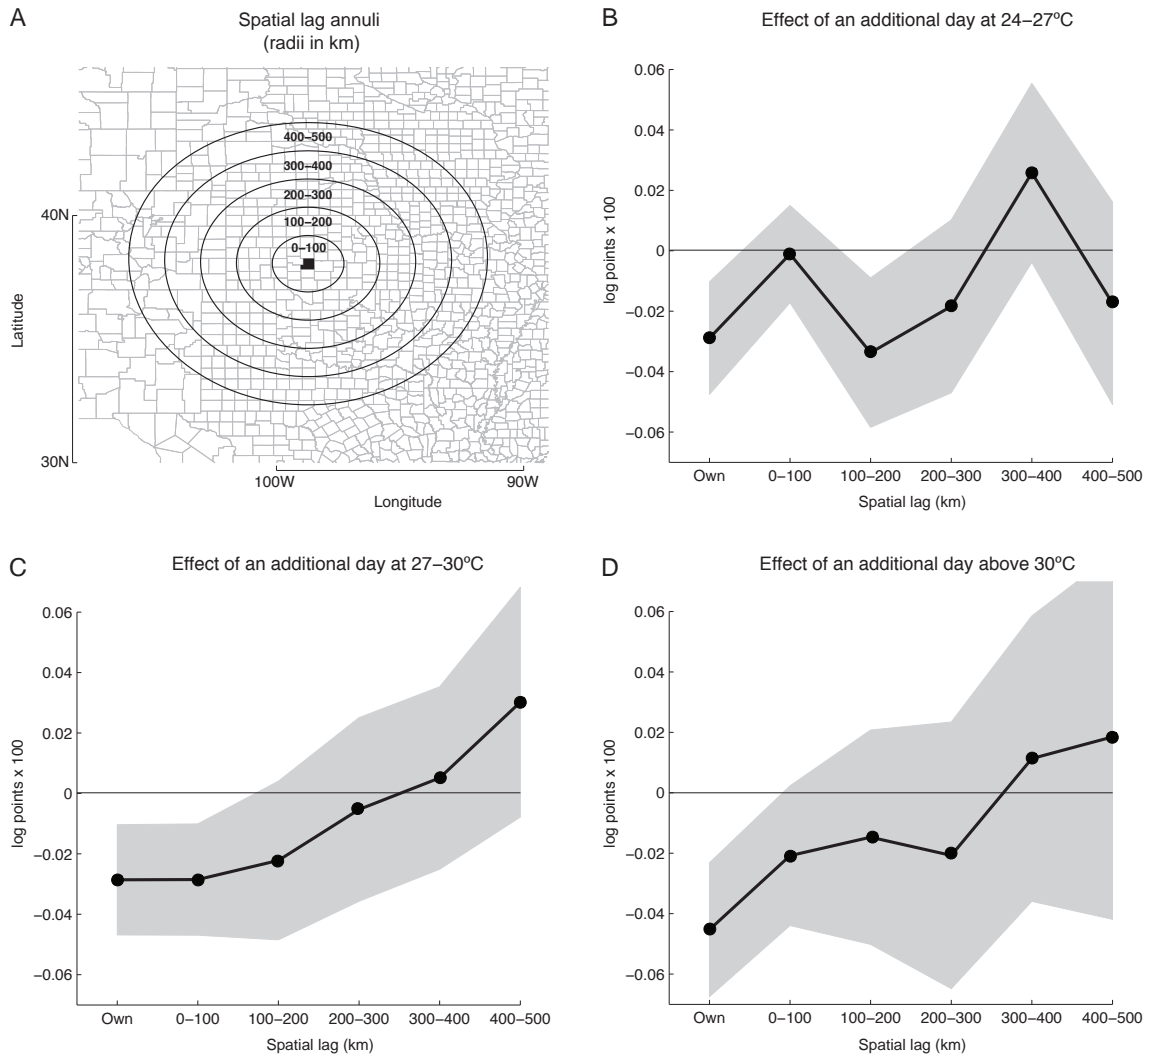


Figure A5: (A) Example of annuli used to construct spatial lags, relative to Stafford, Kansas (black). (B) Effect on i of each additional day at 24–27°C for i and 24–27°C days experienced by j 's at various distances from i . (C) Same but for 27–30°C. (D) Same but for > 30°C. All effects in (B)–(D) are estimated simultaneously, along with own effects for lower temperatures and all controls.




Figure A6: Left: 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. Right: 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). Compare to Figure 9 in the main text.

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 Schlenker and Roberts (2009) study), 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 is consistent with the hypothesis that high-temperature yield declines cause a decline in income that is not offset by rising prices.

The slightly higher breakpoint of 29-32°C in Schlenker and Roberts (2009) and its steeper decline is likely because the authors 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 9 as consistent with the crop yield response reported by Schlenker and Roberts (2009).

Non-Farm Income Our estimated effect of temperature on non-farm income is roughly four times larger than 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. According to Figure 9 in the main text, a day with an average temperature of 25°C causes annual non-farm

income to fall by 0.000213 log points, which corresponds to a loss of 7.8% of an average 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. (2015)), 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.

As with the crop yield response, the breakpoint documented by Graff Zivin and Neidell (2014) ($\sim 25^\circ\text{C}$) 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.³⁹ 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) and Parsons (2014) indicate that productivity begins to decline at slightly lower temperatures ($\sim 21\text{-}22^\circ\text{C}$). Finally, we also 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 estimate at the adjacent temperature bin.

A6 Point estimates for key results

The table below present select point estimates for our main models (columns 1-3), for a model where the dependent variable is log earnings per capita (column 4) and for a model where we do not apply the log transformation to total income per capita (column 5). All other estimates are available upon request.

³⁹The 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.

Table A2: The effect of temperature on income

	(1)	(2)	(3)	(4)	(5)
	Personal income (log), ×100	Farm income	Non-farm earnings (log)×100	Earnings (log)×100	Personal income
above 30°C	-0.076*** (0.013)	-21.01*** (3.14)	-0.0009 (0.0104)	-0.123*** (0.020)	-22.19*** (3.47)
27 to 30°C	-0.065*** (0.011)	-15.34*** (2.40)	-0.0159** (0.0065)	-0.109*** (0.017)	-18.46*** (2.88)
24 to 27°C	-0.058*** (0.010)	-11.82*** (2.30)	-0.0213*** (0.0059)	-0.087*** (0.016)	-15.65*** (2.71)
21 to 24°C	-0.029*** (0.009)	-6.43*** (1.99)	-0.0119** (0.0054)	-0.039*** (0.014)	-7.54*** (2.38)
18 to 21°C	-0.031*** (0.009)	-5.77*** (1.99)	-0.0118** (0.0053)	-0.038*** (0.014)	-8.40*** (2.48)
15 to 18°C	-0.011 (0.008)	-2.06 (1.75)	-0.0041 (0.0052)	-0.007 (0.013)	-3.58* (2.12)
9 to 12°C	-0.004 (0.008)	-1.47 (1.75)	-0.0086* (0.0051)	-0.002 (0.012)	-0.92 (2.19)
6 to 9°C	-0.000 (0.008)	-0.17 (1.70)	-0.0094* (0.0053)	0.008 (0.012)	-1.66 (2.25)
3 to 6°C	-0.007 (0.009)	-1.86 (1.95)	0.0010 (0.0057)	-0.009 (0.013)	-3.12 (2.46)
0 to 3°C	-0.012 (0.009)	-3.18 (1.98)	-0.0042 (0.0060)	-0.018 (0.015)	-3.93 (2.60)
-3 to 0°C	-0.020* (0.010)	-4.44** (2.24)	-0.0174*** (0.0065)	-0.039** (0.016)	-5.50** (2.79)
-6 to -3°C	-0.032*** (0.012)	-6.58** (2.65)	0.0003 (0.0079)	-0.035* (0.018)	-10.42*** (3.28)
-9 to -6°C	-0.027** (0.012)	-3.90 (2.91)	-0.0010 (0.0088)	-0.019 (0.020)	-7.86** (3.55)
-12 to -9°C	-0.014 (0.017)	-2.48 (3.61)	-0.0108 (0.0100)	-0.024 (0.026)	-2.59 (4.62)
-15 to -12°C	0.013 (0.027)	3.50 (6.38)	0.0172 (0.0140)	0.001 (0.044)	5.17 (7.32)
below -15°C	0.023 (0.027)	8.59 (6.36)	-0.0112 (0.0096)	0.038 (0.042)	7.10 (7.01)
Observations	76,576	76,576	76,576	76,573	76,576
R-squared	0.587	0.26	0.8534	0.618	0.70

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, with columns indicating whether a log transformation was applied. Controls include year and county fixed effects, lagged weather variables and the lagged dependent variable. Omitted category is 12-15°C.