The Marginal Product of Climate*

Tatyana Deryugina
University of Illinois,
Urbana-Champaign
NBER

Solomon Hsiang
University of California,
Berkeley
NBER

(Preliminary)

Abstract

We develop an empirical approach to value marginal changes to a climate in terms of total market output given optimal factor allocations in general equilibrium. Our approach fully accounts for unobservable heterogeneity as well as all costs and benefits of adaptation in climates of arbitrary dimension. Importantly, we use the Envelope Theorem to show that the marginal product of a probabilistic long-run climate can be exactly identified using only idiosyncratic weather realizations. We apply this approach to the temperature climate of the modern United States and find that, despite evidence that populations have adapted to their local climates, the marginal product of climate has remained unchanged during 1970-2010, with the highest temperatures having the lowest net value. Capital investments associated with urbanization are important but incomplete substitutes for mild temperatures. Integration of marginal products allows us to construct a value function for climate up to a constant, allowing valid causal, non-marginal, cross-sectional and climate change comparisons net of all re-optimization using existing adaptation technologies. In our preferred specification, we estimate that, for example, the climate of Northern Minnesota returns over $2,000 per capita more annually than the climate of Southern Texas. Using a 3% discount rate, the NPV of "business as usual" warming (RCP8.5) until 2100 in the median scenario is a loss of $6.7 trillion.

*Portions of this analysis were previously circulated as a 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 Deschênes, Don Fullerton, Michael Greenstone, Hilary Hoynes, Amir Jina, Bentley Macleod, Kyle Meng, Edward Miguel, Robert Pindyck, Billy Pizer, Julian Reif, James Rising, Michael Roberts, Wolfram Schlenker, Richard Schmalensee, Joe Shapiro, Tony Smith, James Stock, Anna Tompsett, Reed Walker, Frank Wolak, and seminar participants at Chicago, IZA, LSE, MIT, Princeton, Stanford, UC Berkeley, UC Davis, UC Santa Barbara, UI Urbana-Champaign, UNC, and Yale for discussions and comments. We thank DJ Rasmussen, Michael Delgado, and Andrew Wilson for research assistance. This version: July 23, 2017

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 derive how our approach can recover the marginal product of climate by only exploiting idiosyncratic weather variation using a reduced-form estimator.

Our core insight is that the marginal effect of weather, suitably defined, is exactly equal to the marginal effect of climate on total output in the "local neighborhood" of a long-run market equilibrium—where neighborhood is defined in the space of all possible probability distributions of environmental states, which we term the *space of all possible climates*. As is well understood, the short-run income response to weather and the long-run income response to climate are not the same mathematical object, due to possibility of adaptive re-optimization in the long-run (Mendelsohn, Nordhaus, and Shaw, 1994; Kelly, Kolstad, and Mitchell, 2005). However, the two surfaces that describe these relationships are exactly tangent in any observed equilibrium, which occurs because no adaptive reallocation occurs in the short-run, by definition, and the effect of marginal adaptation on income in the long-run is exactly zero, a direct result of the Envelope Theorem (Guo and Costello, 2013). This equivalence enables us to measure the local gradient of the long-run relationship implicitly by measuring the local gradient of the short-run relationship directly. Recognizing that the two local gradients are identical is useful because the short-run response can be empirically identified by exploiting idiosyncratic variations in weather over time within each location, allowing the econometrician to purge unobserved cross-sectional heterogeneity from these parameter estimates. 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, (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 for a large number of 'small macro-economies' represented by the panel of modern US counties, which plausibly satisfy the criteria above, when we examine their local response to small perturbations in the annual distribution of daily temperatures. We discover a remarkably strong and stable relationship between temperatures and production across space, seasons, and over time. Importantly, we show that large 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 into 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 non-marginal effects on production due to large climate changes. Thus, we can compute differences in economic production that are attributable to differences in contemporaneous locations' climates, as well as to estimate distortions to future production due to projected future warming. In both cases, our construction of this integral accounts for both the costs and benefits of all margins of endogenous adaptation along the entire evolution of a high dimensional climate. On net, we find that existing climate differences between counties generates substantial differences in output, with hotter climates having lower production on average *ceteris paribus*, a result that we recover without exploiting cross-sectional variation in our parameter estimates.

For similar reasons, projections of future output are substantially reduced once future warming and resulting adaptations are both accounted for. Net of all currently available adaptation technologies, we value the projected change in US production during the twenty-first century at $6.7 trillion in NPV (RCP 8.5, discounted at 3% annually) in the median scenario, using our preferred specification. Accounting for climate model uncertainty (Burke et al., 2015), the 90% confidence interval of this estimate is $4.7-10.4 trillion. Importantly, these values do not represent welfare calculations, they do not account for non-market impacts of warming, they do not account for effects of climate change other than temperature changes, and they clearly cannot account for possible future technological innovations that do not yet exist in our data. Interestingly, we note that accounting for adaptive reallocations using our approach *increases* total projected losses relative to a naive approach that assumes uniform marginal effects everywhere, a result that is counter to the widely referenced "folk theorem" that the latter should be larger in magnitude. This occurs because 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, in models that assume uniform marginal effects, we under-estimate total future losses because parameter estimates for northern locations are biased toward zero by pooling those counties with hotter and more adapted (but less productive) counties in the South, thereby obscuring the differentially larger effects in the North.

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. Our definitions are somewhat more formal than is standard in this literature, but as we show, construction of this new formulation is much of what delivers our main theoretical result. 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 nonmarginal effects of climate. We also provide graphical illustrations to build intuition for this result in simple climates of one and two dimensions. In Section 4 we explain our empirical implementation in the modern US, deriving how our empirical specification recovers 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 Statement of the problem

To fully quantify the 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 any adaptation measures taken in response. Because the climate is a joint probability distribution over a large number of possible environmental conditions that may occur at a given location and time, we necessarily are looking for a mapping that takes us from the space of possible joint probability distribution functions to a scalar measure of production. Since this input to the economy that we are focusing on is a function, the transformation of climate and other inputs into economic output (a scalar) will be governed by a *functional*[1]. The core challenge of our analysis is to find a suitable framework for translating this functional into an empirically tractable object while simultaneously accounting for endogenous adaptation to any changes in climate.

2.1 Defining climate and its marginal product

Let the relevant environmental conditions, which we consider state variables, at location i and time t be represented by the vector $\mathbf{x}$:

$$\mathbf{x}_{it} = [temperature_{it},\ precipitation_{it},\ humidity_{it}, ...]\,, \tag{1}$$

where $\mathbf{x}_{it}$ is a draw from the joint probability distribution function $f(\mathbf{x})$. We are interested in how changes to this probability distribution function alter the economic value of output in a given economy. The functional $Y(.)$ maps this function to economic output:

$$f(\mathbf{x}) \mapsto Y(f(\mathbf{x}), \mathbf{b}), \tag{2}$$

where $\mathbf{b}$ is the vector of length N describing all endogenous control variables in the economy, which must also include all inputs not described by $\mathbf{x}$. Alteration of these control variables is how the economy may adapt to changes in $f(\mathbf{x})$.

[1]A functional is similar to a function, but takes functions (rather than scalars or vectors) as arguments and outputs a scalar. Thus, functionals are a mapping from a function space to the space of reals. A definite integral is a commonly used functional.

Define Ψ to be the function space of valid probability distribution functions over the state variables in $\mathbf{x}$:

$$f(\mathbf{x}) \in \Psi. \tag{3}$$

The structure of Ψ is constrained both by physics and the identity that the joint distribution and its marginal distributions all integrate to one.

To make the characterization of (2) empirically tractable, we exploit a "function-generating function" $\psi(.)$ that describes the entire probability distribution function $f(\mathbf{x})$ in terms of a vector $\mathbf{C}$:

$$\mathbf{C} \mapsto \psi(\mathbf{C}, \mathbf{x}) = f(\mathbf{x}). \tag{4}$$

Output from this function-generating function must also be a valid probability distribution:

$$\psi(\mathbf{C}, \mathbf{x}) \in \Psi. \tag{5}$$

Let the vector $\mathbf{C}$ have length K and lie in the space $\mathcal{C}$. $\mathcal{C}$ is defined such that

$$\mathbf{C} \in \mathcal{C} \iff \psi(\mathbf{C}, \mathbf{x}) \in \Psi. \tag{6}$$

Thus, for any vector $\mathbf{C}$ in the space $\mathcal{C}$, there exists a valid probability distribution over $\mathbf{x}$ generated by $\psi(\mathbf{C}, \mathbf{x})$. Thus, without loss of generality, we can rewrite $Y(f(\mathbf{x}), \mathbf{b})$ as $Y(\mathbf{C}, \mathbf{b})$, where it is understood that $\mathbf{C}$ is the parameterization of the probability distribution $f(\mathbf{x})$ through the function $\psi(.)$. The purpose of using $\psi(.)$ in this formulation is that it creates a parameterized mapping from the function space Ψ to the vector space $\mathcal{C}$ such that we can treat the functional $Y(.)$ as if it were a function, so long as we know the mapping function $\psi(.)$.

By construction, the vector $\mathbf{C}$ is a set of sufficient statistics for $f(\mathbf{x})$. For simplicity, we will refer to this vector as *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

$$\frac{\mathrm{d}Y(\mathbf{C}_0, \mathbf{b}_0^*)}{\mathrm{d}\mathbf{C}} = \lim_{\mathbf{C}' \to 0} \left[Y(\mathbf{C}_0 + \mathbf{C}', \mathbf{b}_0^* + \mathbf{b}^{*\prime}) - Y(\mathbf{C}_0, \mathbf{b}_0^*) \right], \tag{7}$$

where the vectors $\mathbf{b}_0^*$ and $\mathbf{b}_0^* + \mathbf{b}^{*\prime}$ are endogenous to the economy and, presumably, optimized to adapt the economy to the two climates $\mathbf{C}_0$ and $\mathbf{C}_0 + \mathbf{C}'$.

2.2 Information and adaptation to climate in general equilibrium

In general equilibrium, individuals with knowledge of the state $\mathbf{C}$ will adjust their behavior and factor allocations to maximize their private utility. If each individual behaves as a price-taker and prices are non-zero, these decentralized choices results in maximization of the value of total revenue in the market economy Y (Arrow and Debreu, 1954; Koopmans, 1957). The vector $\mathbf{b}$ describes all N control variables throughout the economy, including all production, consumption,

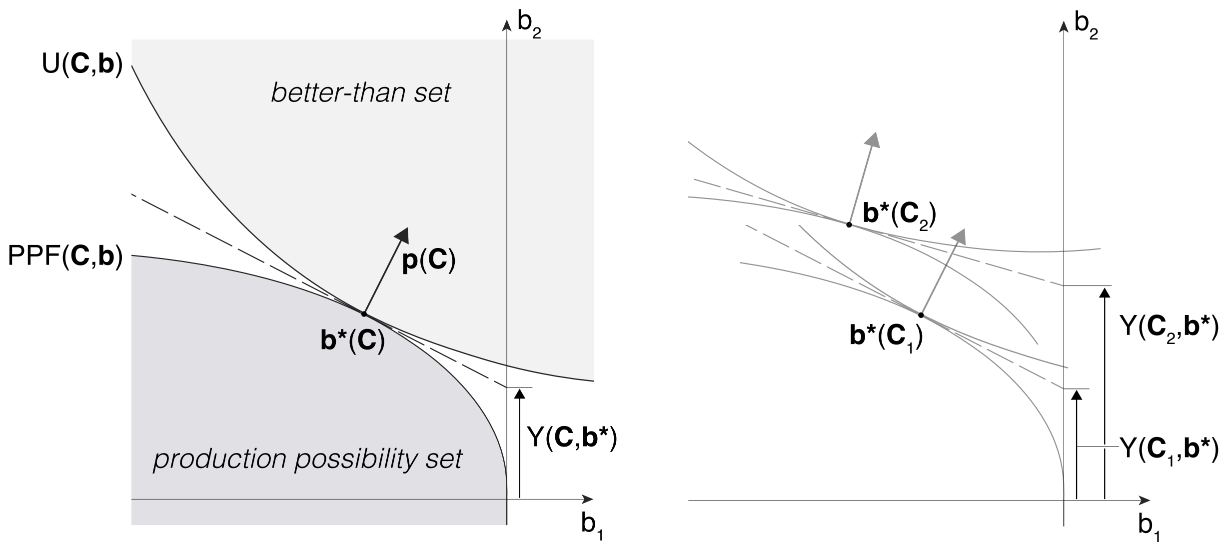

Figure 1: General equilibrium in a climate $\mathbf{C}$ where $\mathbf{b}_1$ (an input to production as drawn) and $\mathbf{b}_2$ (an output) are endogenously determined. Left: Preferences $U(.)$ and the production possibility set may both be influenced by climate. The equilibrium allocation is $\mathbf{b}^*(\mathbf{C})$ with price vector $\mathbf{p}$. Total market revenue (income) is Y, treating b_2 as the numeraire. Right: As the climate of an economy is altered from $\mathbf{C}_1 \to \mathbf{C}_2$, both the PPF and the better-than set will adjust, producing a new optimal allocation $\mathbf{b}^*(\mathbf{C}_1) \to \mathbf{b}^*(\mathbf{C}_2)$, an adjustment of the price vector, and a change in total revenue Y.

and investment decisions. In response to the information contained in $\mathbf{C}$, $\mathbf{b}$ will be optimally set to $\mathbf{b}^*$ in equilibrium, leading us to write $\mathbf{b} = \mathbf{b}(\mathbf{C})$. Thus, in an economy with full information of the climate and N margins of adaptation, $\mathbf{b}$ is endogenously set to

$$\mathbf{b}^*(\mathbf{C}) = \arg\max_{\mathbf{b}} Y(\mathbf{C}, \mathbf{b}(\mathbf{C})). \tag{8}$$

Thus it is clear that, as far as total revenue in the market 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 were to induce an endogenous adjustment in $\mathbf{b}^*$, then it must be the case that any reductions in Y caused by this adjustment were outweighed by its gains under the new climate. To illustrate this, consider an economy in general equilibrium where $N = 2$, as shown in the left panel of Figure 1, where representative utility is $U(\mathbf{C}, \mathbf{b})$ and the production possibility frontier is $PPF(\mathbf{C}, \mathbf{b})$. If the climate $\mathbf{C}$ changes then the economy will adapt by reallocating factors $\mathbf{b}_1$ and $\mathbf{b}_2$ to maximize $U(.)$, which in turn will necessarily maximize total revenue in the economy, i.e. income Y, given nonzero prices $\mathbf{p}(\mathbf{C})$ (Koopmans, 1957).

We note that if some factors have no market price, such as a completely externalized pollutant or a non-market good, they will not be factored into the market and will not affect total revenue for the market. This is not an immediate concern for us because we are focused on changes to total market output, but it is worth noting because many other approaches to valuing climatic conditions in welfare terms do take into account both market and non-market responses to climate (e.g. Anthoff, Hepburn, and Tol (2009); Hsiang et al. (2017)).

This formulation of adaptation is general in the sense that it accounts for an arbitrary number of endogenous adjustments to climate, including their opportunity costs, 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—both of which may also have opportunity costs. Clearly, in all cases prices will adjust, but it is not required that we observe them directly in order to measure the marginal product of climate. To construct this marginal product, we simply must observe how Y responds to changes in $\mathbf{C}$, net of all general equilibrium adjustments, as shown in the right panel of Figure 1.

Because our focus is the value of the climate in terms of total market production net of all optimal re-adjustments, we can write equilibrium output as the value function

$$V(\mathbf{C}) = Y(\mathbf{C}, \mathbf{b}^*(\mathbf{C})), \tag{9}$$

which captures the net costs and benefits of all possible adaptations embodied by $\mathbf{b}^*$. If there exist regions in the climate space $\mathcal{C}$ where endogenous adjustment of $\mathbf{b}$ can fully offset any changes in total production induced by changes in $\mathbf{C}$, than the value function $V(\mathbf{C})$ will be flat in that region. However, if compensating adjustments in $\mathbf{b}$ are not possible or impose high costs (benefits) by moving resources away from (towards) other productive activities, then the value function may have a steeper gradient.

Throughout this analysis we assume that $\mathbf{b}$ lies in the state space $\mathcal{B}$ which is a dense subset of $\mathbb{R}^N$; further we assume that Y is concave and differentiable in the elements of $\mathbf{b}$. These assumptions ensure that $\mathbf{b}^*$ exists. In an empirical application, this implies that the elements of $\mathbf{b}$ take on continuous values, a natural assumption for most standard variables in an economy of sufficiently large scale, such as the number of apples are sold in a regional market or the number of miles of road laid down within a county. For some specific control variables often discussed as margins of adaptation to climate, such as the construction of sea-walls or the switching of crops that are planted, the literature often frames these decisions as discrete.[2] However, such a framing is not generally accurate even at the scale of a decision-maker because many decisions have (continuous) intensive margins. For example, a sea-wall can always be slightly longer or slightly higher and a farmer can always allocate just a slightly larger fraction of cropland to a new variety. At larger scales of aggregation—such as at county, state, or national levels—the assumption that elements

[2]For example, both Mendelsohn, Nordhaus, and Shaw (1994) and Deschênes and Greenstone (2007) describe crop-switching as a discrete choice variable.

of $\mathbf{b}$ takes on continuous values is even more easily defended, as the economic decisions of many individuals are easily aggregated into continuous measures.

2.3 Weather and climate

Once agents use their knowledge of $\mathbf{C}$ to make all behavioral and factor allocation decisions, described by $\mathbf{b}$, there remain no possible channels for the abstract probability distribution described by $\mathbf{C}$ to influence economic outcomes except through generating actual realizations of $\mathbf{x}_{it}$. A climate that is generally wet will generate more rainy days and a climate that is generally warm will generate more hot days. These actual events will potentially affect Y, perhaps mediated by adaptations described by $\mathbf{b}^*(\mathbf{C})$.

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

$$\mathbf{c}_{i\tau} \mapsto \psi(\mathbf{c}_{i\tau}, \mathbf{x}) = \hat{f}(\mathbf{x})_{i\tau}. \tag{10}$$

By construction, $\mathbf{c}_{i\tau}$ has the same dimensionality as $\mathbf{C}$ and also lies in $\mathcal{C}$. We define $\mathbf{c}_{i\tau}$ as the *weather* realized at i during period τ.

The description of $\mathbf{c}_{i\tau}$ may appear on its face more complicated than vernacular usage of the term "weather," although we believe it actually maps very closely to common usages of the term. It is tempting to consider $\mathbf{x}_{it}$ the weather, since these are directly observable realizations, but we caution that elements in $\mathbf{x}_{it}$ are continuous measures taken at continuous (or at least infinitesimally short) moments in time indexed by t. In contrast, $\mathbf{c}_{i\tau}$ summarizes many of these measures taken over a finite interval of time. Thus, if one were to ask "what was today's weather?" a reasonable response would be "pretty warm," or "a high of 80 and low of 60," both of which are summary statements more akin to the summary captured in $\mathbf{c}_{i\tau}$ than to an infinite number of measures of temperature taken at all moments t within a day, as is captured by $\mathbf{x}_{it}$.

Because C summarizes the probability distribution function $f(\mathbf{x})$, which produces realizations $x_{i\tau}$ that are in turn is used to construct the summary $c_{i\tau}$, it is straightforward to consider $c_{i\tau}$'s as random vectors generated by some function of C. Stated another way, weather is a random realization of events that are determined by the climate. To capture this intuition and simplify notation, we write

$$\mathbf{c} = \mathbf{c}(\mathbf{C}) \tag{11}$$

to denote some distribution of weather realizations $\mathbf{c}_{i\tau}$ realized from a climate $\mathbf{C}$.

Because the probability distribution described by the climate $\mathbf{C}$ can only affect economic production (1) as information, through its effect on decisions embodied by $\mathbf{b}^*$, and (2) through its influence on actual events, through its effect on weather realizations $\mathbf{c}$, it is useful to rewrite to Eq.

9 so that $\mathbf{C}$ clearly enters only through these two arguments:

$$Y(\mathbf{C}, \mathbf{b}(\mathbf{C})) = Y(\mathbf{c}(\mathbf{C}), \mathbf{b}(\mathbf{C})). \tag{12}$$

Using this notation we write down our final formulation of the value function for climate as the optimal endogenous output by an economy when weather events are realized from a probabilistic climate and all factors and production decisions are optimized with full knowledge of the climate

$$V(\mathbf{C}) = Y(\mathbf{c}(\mathbf{C}), \mathbf{b}^*(\mathbf{C})). \tag{13}$$

The shape of the income-generating function $Y(.)$ along the optimized path $\mathbf{b}^*(\mathbf{C})$ over all $\mathbf{C} \in \mathcal{C}$ ultimately determines the value of the climate $V(\mathbf{C})$, net of all adaptation costs and benefits captured by the market. Thus, the full marginal product of the climate from Eq. 7 can be rewritten as the local gradient in the value function

$$\frac{\mathrm{d}Y(\mathbf{c}(\mathbf{C}), \mathbf{b}^*(\mathbf{C}))}{\mathrm{d}\mathbf{C}} = \frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}} = \lim_{\mathbf{C}' \to 0} \left[V(\mathbf{C} + \mathbf{C}') - V(\mathbf{C})\right], \tag{14}$$

where $\mathbf{C}'$ describes the structure of an arbitrary perturbation to the current climate vector $\mathbf{C}$ such that $\mathbf{C} + \mathbf{C}' \in \mathcal{C}$. Empirically recovering the gradient vector $\frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}}$ for the modern counties in the United States is the goal of this analysis.

2.4 Relationship of this formulation to prior work

Seminal analysis by Mendelsohn, Nordhaus, and Shaw (1994) attempted to directly estimate Equation 13 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 $\psi(.)$, 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 $\psi(.)$, 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 $\hat{\mathbf{C}} = \mathbf{c}_{i\tau}$ and then estimate a version of Equation 13 exploiting random variation in $\mathbf{c}_{i\tau}$. 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 $\mathbf{c}_{i\tau}$ 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 $\psi(.)$ 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. (2013), 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 $\mathbf{c}_{i\tau}$ on Y in Equation 13 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 τ when constructing $\hat{\mathbf{C}} = \mathbf{c}_{i\tau}$, 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 $\psi(.)$, allowing for nonlinear adaptation at all points in the distribution of $f(\mathbf{x})$, and restricting our analysis to an optimized outcome where short-run marginal changes in $\mathbf{c}_{i\tau}$ exactly identify the marginal effect of long-run changes in $\mathbf{C}$.

3 Identifying the full marginal product of climate empirically

As formulated here, an economy's output $Y(\mathbf{C}, \mathbf{b})$ depends on its position in the $K+N$ dimensional Cartesian space $\mathcal{C} \times \mathcal{B}$, recalling that $\mathbf{C} \in \mathcal{C}$, $\mathbf{c} \in \mathcal{C}$, and $\mathbf{b} \in \mathcal{B}$. Exogenous changes in the position of the economy in the subspace of $\mathbf{C}$ lead to endogenous re-optimization of state variables $\mathbf{b}$ to $\mathbf{b}^*$ such that income $Y(\mathbf{c}(\mathbf{C}), \mathbf{b}^*(\mathbf{C}))$ is always at an optimum with respect to the subspace of $\mathbf{b}$. A core empirical challenge has been tracing out the optimal path of an economy through the space $\mathcal{C} \times \mathcal{B}$ as it adjusts $\mathbf{b}^*$ in response to changes in position $\mathbf{C}$.

Our objective is to characterize the marginal product of climate

$$\frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}} = \left[\frac{\partial V}{\partial \mathbf{C}_1}, \quad \cdots, \quad \frac{\partial V}{\partial \mathbf{C}_K}\right]\bigg|_{\mathbf{C}} \tag{15}$$

where $\mathbf{C}_k$ denotes the k-th element in $\mathbf{C}$. Note that the marginal product of climate is equivalent to the gradient of the value function $\nabla V(\mathbf{C})$ in the space of $\mathcal{C}$, so as a vector it points in the direction of most rapid increase in the value function everywhere. Since the value function $V(\mathbf{C})$ is the outcome $Y(.)$ when $\mathbf{b}$ is optimized, we decompose $\frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}}$ into contributions that come from (1) the direct effect of changing $\mathbf{C}$ and the resulting changes in $\mathbf{c}$ and (2) the endogenous adjustments to $\mathbf{b}$ in response to the knowledge that $\mathbf{C}$ has changed. Following Hsiang (2016), we term these contributions the *direct effect* and the *belief effect*—since the latter is entirely driven by changes in beliefs of individuals over the structure of the unobservable probability distribution $f(\mathbf{x})$, and thus implicitly by their beliefs over $\mathbf{C}$. We write

$$\frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}} = \frac{\mathrm{d}Y(\mathbf{c}(\mathbf{C}), \mathbf{b}^*(\mathbf{C}))}{\mathrm{d}\mathbf{C}} = \underbrace{\sum_{k=1}^{K} \frac{\partial Y(\mathbf{C})}{\partial \mathbf{c}_k} \frac{\mathrm{d}\mathbf{c}_k}{\mathrm{d}\mathbf{C}}}_{\text{direct effects}} + \underbrace{\sum_{n=1}^{N} \frac{\partial Y(\mathbf{C})}{\partial \mathbf{b}_n} \frac{\mathrm{d}\mathbf{b}_n^*}{\mathrm{d}\mathbf{C}}}_{\text{belief effects}} \tag{16}$$

where $\frac{\mathrm{d}\mathbf{c}_k}{\mathrm{d}\mathbf{C}}$ and $\frac{\mathrm{d}\mathbf{b}_n}{\mathrm{d}\mathbf{C}}$ are the k-th and n-th row vectors of Jacobians $\frac{\mathrm{d}\mathbf{c}}{\mathrm{d}\mathbf{C}}$ (size $K \times K$) and $\frac{\mathrm{d}\mathbf{b}_n}{\mathrm{d}\mathbf{C}}$ (size $N \times K$), respectively.[3] Note that each partial derivative is evaluated "locally" in the neighborhood of the initial climate $\mathbf{C}$.

In the ideal experiment, we would compare two identical economies with identical climates $\mathbf{C}$. We would then perturb the climate of one by the small disturbance $\mathbf{C}'$ to $\mathbf{C} + \mathbf{C}'$ and allow the economy to adjust $\mathbf{b}^*$ endogenously. The difference in productivities between the treatment and control economies would then be the marginal product $\frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}}$, capturing contributions from both direct effects and belief effects.

Previous work, such as Mendelsohn, Nordhaus, and Shaw (1994), approximated this experiment in a cross-sectional hedonic framework, where productivities were conditioned on climate $\mathbf{C}$ and covariates. The benefit of this approach is that it captures belief effects that differ between contemporaneous counties. However, a drawback of this approach is that it cannot ensure that unobservable heterogeneity in determinants of productivity is fully accounted for.

Later work, such as Schlenker and Roberts (2009) and Deschênes and Greenstone (2011), exploited exogenous variation in weather $\mathbf{c}$ within a location to purge estimates of any omitted variables bias due to unobservable heterogeneity. This approach identifies direct effects, but has been critiqued for failing to capture belief effects, since beliefs about the climate may not respond to weather ($\frac{\partial \mathbf{b}^*}{\partial \mathbf{c}} \neq \frac{\partial \mathbf{b}^*}{\partial \mathbf{C}}$).

Here we derive conditions under which the marginal product of climate, as described by Eq 16, is exactly identified using a within-location estimator that exploits random variation in weather. This approach builds on the methods proposed by Schlenker and Roberts (2009) and Deschênes and Greenstone (2011) to purge estimates of omitted variables bias from unobservable heterogeneity, but requires more stringent conditions that are not met by these previous analyses. Furthermore,

[3] The Jacobian matrices are $\frac{\mathrm{d}\mathbf{c}}{\mathrm{d}\mathbf{C}} = \begin{bmatrix} \frac{\partial \mathbf{c}_1}{\partial \mathbf{C}_1} & \cdots & \frac{\partial \mathbf{c}_1}{\partial \mathbf{C}_K} \\ \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{c}_K}{\partial \mathbf{C}_1} & \cdots & \frac{\partial \mathbf{c}_K}{\partial \mathbf{C}_K} \end{bmatrix}$ and $\frac{\mathrm{d}\mathbf{b}}{\mathrm{d}\mathbf{C}} = \begin{bmatrix} \frac{\partial \mathbf{b}_1}{\partial \mathbf{C}_1} & \cdots & \frac{\partial \mathbf{b}_1}{\partial \mathbf{C}_K} \\ \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{b}_N}{\partial \mathbf{C}_1} & \cdots & \frac{\partial \mathbf{b}_N}{\partial \mathbf{C}_K} \end{bmatrix}$.

integration of marginal products of climate while accounting for belief effects (i.e. adaptation), which is necessary to compute the effects of non-marginal changes in climate, requires a new method that we derive.

3.1 Empirically identifying effects of climate using weather

Variation in climate $\mathbf{C}$ and variation in weather $\mathbf{c}$ are not the same. Further, the marginal product of climate ($\frac{dV(\mathbf{C})}{d\mathbf{C}}$) and the marginal product of weather ($\frac{dY(\mathbf{C})}{d\mathbf{c}}$) are different mathematical objects. However, we show here that the values of these two marginal products are equal under certain conditions, allowing us to measure the value of climate by isolating the marginal product of weather net of all unobservable heterogeneity, which is generally easier than isolating the marginal product of climate directly. This insight results directly from application of the Envelope Theorem.[4]

The intuition of the result is as follows. Imagine there are two otherwise identical households that are next-door neighbors on a street that runs North-South. The more northern household faces a *very slightly* different climate because it is *very slightly* further north. The difference in climate faced by the two households is vanishingly small, but nonzero. These two households have the ability to adapt many dimensions of their daily life to their beliefs about their respective climates and will adopt slightly different behaviors and investments that maximize various outcomes, generating belief effects. However, if we focus on outcomes that are maximized by the households, then the overall net effect caused by these slightly different adaptation decisions is zero because any marginal benefits that the northern household reaps are exactly offset by additional marginal costs (since the household is at a maximum). Therefore, any difference in the optimized outcome between the two households must come from the direct effects of the slightly different climate, and the influence of slightly different beliefs and adaptations between the two households can be ignored. If a weather realization occurs such that the southern household experiences conditions that are slightly different from what they expect and its distribution of weather actually matches the climate of the northern household, then this "weather effect" on the optimized outcome of the southern household must be exactly the same as the cross-sectional difference across the two households in a year when their weather realizations match their respective climates perfectly—since in both cases there is no influence of changing beliefs on the optimized outcome. Stated simply, the marginal effect of the climate on an optimized outcome is exactly the same as the marginal effect of the weather.

Based on this insight, we can trace out a curve describing climate effects between sequential neighbors by watching how optimized outcomes in each household change when that household is confronted by a weather distribution that matches the climate of their immediate next-door neighbor. The integral of these marginal differences between sequential neighbors must then describe how the climate generates larger differences between households that are not adjacent neighbors

[4]In 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 withe predictions of the Envelope Theorem as it is used in the present analysis.

and experience climates that differ by a non-marginal amount. Importantly, this integration procedure does not assume that individuals do not adapt to their climate. Rather, the marginal effect of such adjustments for marginal climate changes is zero on an optimized outcome, so marginal effects of weather—which do not cause beliefs to change—can be used as a substitute for marginal climate changes in the integration, despite the presence of changing beliefs and adaptations.

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 $\frac{\mathrm{d}V(\mathbf{C}^a)}{\mathrm{d}\mathbf{C}}$ we differentiate V by $\mathbf{C}_k$ (Eq 15). By the chain rule we have

$$\begin{aligned}\frac{\mathrm{d}V(\mathbf{C}^a)}{\mathrm{d}\mathbf{C}_k} &= \frac{\mathrm{d}Y(\mathbf{c}(\mathbf{C}^a), \mathbf{b}^*(\mathbf{C}^a))}{\mathrm{d}\mathbf{C}_k} \\ &= \frac{\partial Y(\mathbf{c}(\mathbf{C}^a), \mathbf{b}^*(\mathbf{C}^a))}{\partial \mathbf{C}_k} + \sum_{\kappa=1}^{K} \frac{\partial Y(\mathbf{c}(\mathbf{C}^a), \mathbf{b}^*(\mathbf{C}^a))}{\partial \mathbf{c}_\kappa} \frac{\mathrm{d}\mathbf{c}_\kappa}{\mathrm{d}\mathbf{C}_k} + \sum_{n=1}^{N} \frac{\partial Y(\mathbf{c}(\mathbf{C}^a), \mathbf{b}^*(\mathbf{C}^a))}{\partial \mathbf{b}_n} \frac{\mathrm{d}\mathbf{b}_n}{\mathrm{d}\mathbf{C}_k}\end{aligned} \tag{17}$$

where

$$\frac{\partial Y}{\partial \mathbf{C}_k} = 0 \tag{18}$$

since 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. Because V is the outcome when Y has been optimized through all possible adaptations, and it is differentiable in $\mathbf{b}$, we also know by the Envelope Theorem that

$$\frac{\partial Y(\mathbf{c}(\mathbf{C}^a), \mathbf{b}^*(\mathbf{C}^a))}{\partial \mathbf{b}_n} = 0 \tag{19}$$

for all N dimensions of the $\mathbf{b}$ subspace. Thus, Eq. 17 simplifies to

$$\begin{aligned}\frac{\mathrm{d}V(\mathbf{C}^a)}{\mathrm{d}\mathbf{C}_k} &= \sum_{\kappa=1}^{K} \frac{\partial Y(\mathbf{c}(\mathbf{C}^a), \mathbf{b}^*(\mathbf{C}^a))}{\partial \mathbf{c}_\kappa} \frac{\mathrm{d}\mathbf{c}_\kappa}{\mathrm{d}\mathbf{C}_k} \\ &= \sum_{\kappa=1}^{K} \frac{\partial V(\mathbf{C}^a)}{\partial \mathbf{c}_\kappa} \frac{\mathrm{d}\mathbf{c}_\kappa}{\mathrm{d}\mathbf{C}_k},\end{aligned} \tag{20}$$

where the second equality holds by the definition of $V(.)$ (Equation 13). Noting 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

$$\frac{\mathrm{d}\mathbf{c}_\kappa}{\mathrm{d}\mathbf{C}_k} = \begin{cases} 1 & \text{for } \kappa = k \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

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 $\mathbf{C}$. This restriction is equivalent to setting $\frac{\mathrm{d}\mathbf{c}}{\mathrm{d}\mathbf{C}}$ equal to the identity matrix[5] (Eq. 21). Then we have

$$\frac{\mathrm{d}V(\mathbf{C}^a)}{\mathrm{d}\mathbf{C}_k} = \frac{\partial V(\mathbf{C}^a)}{\partial \mathbf{c}_k}, \tag{22}$$

which says that the total marginal effect on V of the kth dimension of the climate, evaluated at $\mathbf{C}^a$, is equal to the partial derivative of V with respect to the corresponding dimension of weather, also evaluated at $\mathbf{C}^a$. Locally, the marginal effect of the climate on V is identical to the marginal effect of the weather.[6] Extending this to all k dimensions we have

$$\frac{\mathrm{d}V(\mathbf{C}^a)}{\mathrm{d}\mathbf{C}} = \frac{\partial V(\mathbf{C}^a)}{\partial \mathbf{c}}, \tag{23}$$

stating that the full marginal product of climate is equal to the vector of partial effects of k weather measures, net of all endogenous adaptations. An identical equation can be written for income Y since $Y = V$ in equilibrium. Equation 23 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 Schlenker and Roberts (2009) and Deschênes and Greenstone (2011).

Importantly, for Equation 23 to hold, the outcome V must represent a maximized quantity (Equation 13), which is true in the case of income—studied here—but which may not hold in the case of other outcomes, such as crop yields or mortality risk. This maximization results from individuals considering *both* the costs and benefits of any adaptive adjustment to state variables $\mathbf{b}$, guaranteeing that estimation of Equation 23 empirically captures both of these effects.

The equivalence between marginal effects of climate and weather can be used to construct estimates for non-marginal effects of the climate by integrating marginal effects of weather. For an arbitrary climate $\mathbf{C}^b$, we know that 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 ($\mathcal{C}$) from $\mathbf{C}^a \to \mathbf{C}^b$, starting from $V(\mathbf{C}^a)$:

$$V(\mathbf{C}^b) = V(\mathbf{C}^a) + \int_{\mathbf{C}^a}^{\mathbf{C}^b} \frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}} \cdot \mathrm{d}\mathbf{C}. \tag{24}$$

At each position $\mathbf{C} \in \mathcal{C}$, $\frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}}$ is a vector of differentials describing all the marginal effects of the climate measured "locally" at $\mathbf{C}$ (recall Equation 15). From Equation 23 we know that these differentials with respect to climate can be substituted for using differentials with respect to weather, terms that can be estimated empirically via regression

$$\widehat{\frac{\partial V(\mathbf{C})}{\partial \mathbf{c}}} = \hat{\beta}_{weather}\Big|_{\mathbf{C}} \tag{25}$$

[5] This restriction 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).

[6] Hsiang (2016) calls Eq. 22 the *marginal treatment comparability assumption*, which holds exactly in this case.

where $\hat{\beta}$ is a reduced-form parameter estimate of the marginal effects of weather on total economic production. If $\hat{\beta}$ can be credibly estimated for each value of $\mathbf{C}$ along the path of the line integral, then by combining Equations 23-25 we have the value of a non-marginal change in the climate from $\mathbf{C}^a$ to $\mathbf{C}^b$, in the presence of all adaptation adjustments in $\mathbf{b}$, exploiting only exogenous variation in weather:

$$V(\mathbf{C}^b) - V(\mathbf{C}^a) = \int_{\mathbf{C}^a}^{\mathbf{C}^b} \hat{\beta}_{weather}\Big|_{\mathbf{C}} \cdot \mathrm{d}\mathbf{C}. \tag{26}$$

The total difference in economic output due to a change in the climate is thus computed by integrating a sequence of weather-derived marginal products evaluated at each intermediate value of $\mathbf{C}$.

To summarize, climate and weather are not the same. However, properly formulated, a large number of reduced-form estimates for the local marginal effect of weather on production provides us with enough information to reconstruct the full surface of productivities over the space of possible climates. This result is obtained because, if all control variables are optimized to a baseline climate, additional adaptations to local deviations from that climate have no influence on the marginal effect of climate in a sufficiently small neighborhood (by the Envelope Theorem). Similarly, idiosyncratic weather deviations from the baseline climate also induce no additional adaptations, so the marginal effects of climate and weather on output must be equal. But because non-marginal changes in climate may induce substantial adjustments to control variations, the marginal effects of climate and weather may change, so new marginal effects must be empirically estimated at each position in the K-dimensional climate space. Once enough local marginal effects are estimated throughout the climate space, non-marginal effects of climate can be computed by integrating these marginal effects along any line-integral from the starting climate through the K-dimensional climate space to the final climate.

3.2 Graphical illustration of the solution concept

Some simple examples help clarify our solution to valuing climate. Recall that our goal is to construct an estimate for the value of the position $\mathbf{C}$ in the K-dimensional space of potential climates $\mathcal{C}$, allowing for any possible adjustments in control variables $\mathbf{b}$ based on the position $\mathbf{C}$:

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

We derived that the shape of this surface could be reconstructed, up to a constant of integration, by integrating estimates for the reduced-form marginal effects of weather obtained at each position in $\mathcal{C}$ (Equation 26). Here we graphically illustrate the construction of $V(\mathbf{C})$, and how adaptation is captured, in cases where $K = \{1, 2\}$ and $N = 1$. Our actual empirical implementation for the US has $K = 16$ and unknown N, making it more difficult to visualize.

Example: One dimensional climate with one dimension of adaptation Choice of the form for the function-generating function $\psi(\mathbf{C}, \mathbf{x}) = f(\mathbf{x})$ implicitly determines the structure of the

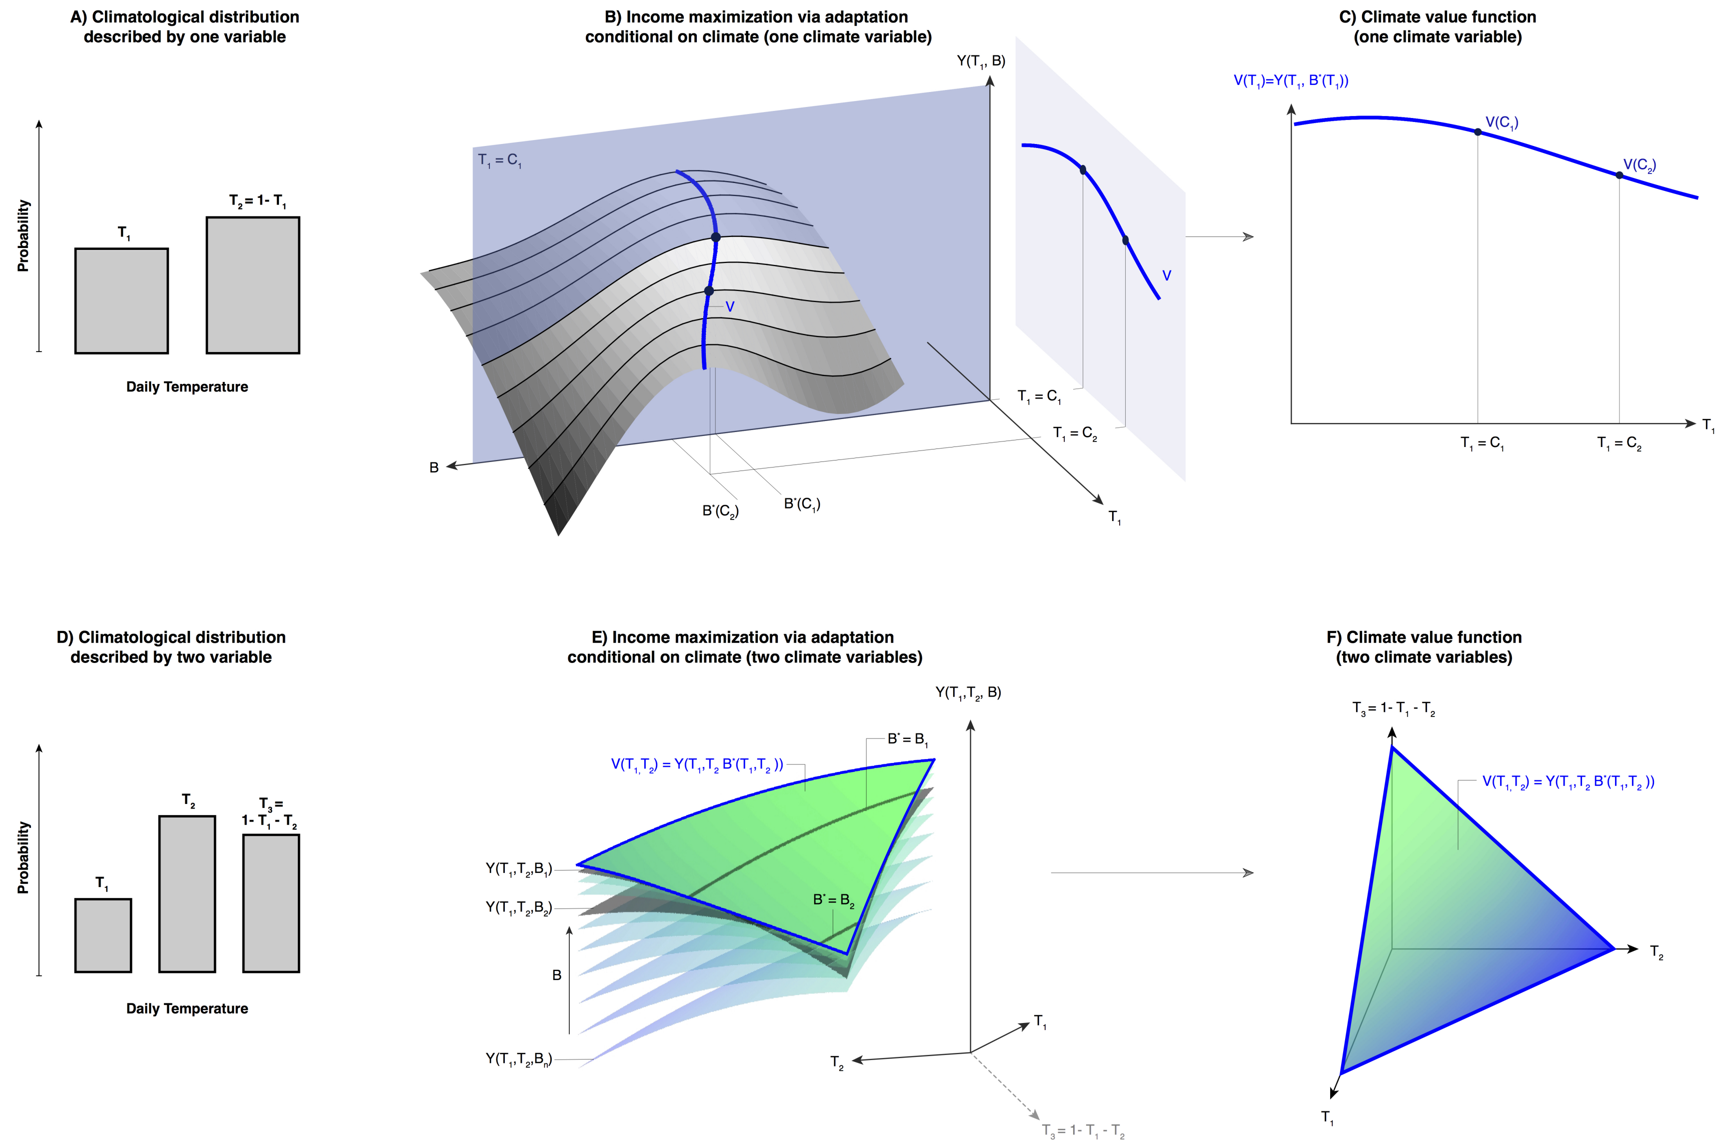

Figure 2: Application of the Envelope Theorem to valuing climate. (A) An example climate defined by a two temperature bin histogram, where the probability distribution of outcomes is fully defined by the mass in one bin (T_1). (B) Height of the gray surface is the income that would be realized for any position in the space defined by the Cartesian product of the climate (T_1, coming out of page) and the control variable (B, positive to left). Conditional on an exogenously determined climate (e.g. blue plane where $T_1 = C_1$), adaptation will seek to maximize output by setting $B = B^*(T_1)$. The blue ridge V represent these optimized quantities. Because these quantities are maximized, small deviations in T_1 from an initial location (black circle) caused by weather induce no adaptation and have slope tangent to V locally. (C) Projection of V onto the T_1 subspace is the value function of this one dimensional climate. (D) Same as (A) except introducing an additional temperature bin. (E) Translucent surfaces are production surfaces over the space (T_1, T_2) for different values of the control variable B. Production maximizes $B = B^*(T_1, T_2)$, such that the opaque triangle (outlined in blue) is the maximum output for each climate position after adaptation, defining the value function $V(T_1, T_2)$. 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). (F) Colors on the 2-simplex depict the value function for each position in the climate space $V(T_1, T_2)$. Note that the height of the simplex surface (T_3) is fully defined by the position (T_1, T_2).

space of all possible climates $\mathcal{C}$, since $\mathbf{C} \in \mathcal{C}$. Begin by considering a one dimensional temperature climate where the probability distribution of daily temperatures is summarized by the scalar T_1, equal to the expected fraction of days in each year below some temperature threshold $\bar{T}$. A second number T_2 equals the expected fraction of days with temperature above $\bar{T}$. This climate is fully described by only one of these values since $T_1 = 1 - T_2$, making this a one dimensional climate. In this example, $\psi(\{T_1\}, \mathbf{x})$ is implicitly defined as the function that places probability mass T_1 below $\bar{T}$ and mass $1 - T_1$ above $\bar{T}$ in a two-bar histogram, as shown in Panel A of Figure 2. Because of the way in which ψ is defined in this example, the space of possible climates is the unit 1-simplex:

$$\mathcal{C} = \{T_1 | T_1 \in [0, 1]\}.$$

The shaded gray surface in Panel B of Figure 2 illustrates the income function $Y(T_1, B)$, which is defined over the space of possible values for T_1 and the single control variable B, which is the only element in $\mathbf{b}$. In a maximizing economy, B is adjusted to maximize Y conditional on the exogenously determined climate T_1. Thus, the maximum income obtainable at any value of T_1 is the "ridge" of the surface[7] that is traced by the blue line. The vertical blue plane at $T_1 = C_1$ illustrates the maximization problem in a given climate: if constrained to select a position on the surface Y along the locations where the $T_1 = C_1$ plane intersects, the economy will adapt to the climate $T_1 = C_1$ by setting the control variable $B = B^*(C_1)$, illustrated by the black circle. If the plane describing the climate were shifted out of the page to $T_1 = C_2$, then the economy would adapt by setting $B = B^*(C_2)$, the second black circle. As the plane describing the climate shifts along all possible values of T_1, the actual trajectory of the adapting economy traces out the blue ridge $Y(T_1, B^*(T_1)) = V$. Projecting V onto the subspace of climate dimensions (light blue plane on right side of Panel B) produces the value function $V(T_1)$ over the space of possible climates, as shown in Panel C.

To illustrate how $V(T_1)$ could be constructed empirically, first consider a reduced form estimate for the marginal effect of weather in this setup for a population with climate T_1. Here, in year t the realized count of days with temperature below $\bar{T}$ would randomly deviate from the expected number T_1 by $\epsilon(t)$. Thus, the one dimension of weather that corresponds to the one dimension of climate is

$$\tilde{T}_1(t) = T_1 + \epsilon(t).$$

Because the control variable B responds to knowledge of the climate but not to random changes in weather, the trajectory of the economy will move along the gray surface in Panel B tracing out the values $Y(T_1 + \epsilon(t), B^*(T_1))$ in the vicinity of a point on the blue ridge. For example, if $T_1 = C_1$, the economy will traverse the surface in the neighborhood of the first black circle as production responds to the small disturbances around C_1. If these disturbances are sufficiently small, then we know from Equation 22 that their tangency will have the same slope, with respect

[7]This ridge is sometimes described as the "envelope" in the Envelope Theorem, although that description is mainly fitting when this three dimensional plot is collapsed to two dimensions (not explicitly showing the B-subspace), as it is usually drawn in microeconomic textbooks.

to T_1, as the curve tracing out $V(T_1)$. This is true—even though the blue line tracing the ridge V moves back and forth in the B-dimension as T_1 changes—because these adjustments occur in a subspace (of B) that is orthogonal to the climate subspace (of T_1) *and* each position on the ridge V is locally optimized with respect to B (Equation 19). Thus, the shape of the blue curve $V(T_1)$ in Panel C can be constructed by integrating a large number of these small tangencies, describing $\frac{\partial Y(\tilde{T}_1, B^*(T_1))}{\partial \tilde{T}_1} = \frac{\partial V(T_1)}{\partial T_1}$, each identified by random variation in weather $\tilde{T}_1(t)$ around each sequential baseline climate T_1.

Example: Two dimensional climate with one dimension of adaptation Expansion of the climate space by one dimension is useful for developing intuition for how this approach generalizes as the dimensionality of the problem increases. Consider restructuring $\psi(.)$ such that the probability distribution of daily temperatures in a year is summarized by a three-bin histogram, shown in Panel D of Figure 2. T_1 is still the expected fraction of days with temperature below $\bar{T}$. However, now T_2 is the expected fraction of days with temperature above $\bar{T}$ and below a second cutoff temperature $\bar{\bar{T}}$. T_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 $T_3 = 1 - T_1 - T_2$. The space of possible climates here is the unit 2-simplex:

$$\mathcal{C} = \{(T_1, T_2) \mid T_1, T_2 \in [0, 1], \;\; T_1 + T_2 \in [0, 1]\}.$$

Let there remain only one dimension of possible adaptive adjustment B.

Depicting $Y((T_1, T_2), B)$ now requires four dimensions. Panel E of Figure 2 depicts multiple semi-translucent surfaces, each a function over the 2-simplex $\mathcal{C}$, holding a value of B_n fixed. The height of the B_n surface at a point (T_1, T_2) is the level of output the economy would exhibit for the climate (T_1, T_2) if $B = B_n$. Optimization of output would lead to selection of $B = B^*(T_1, T_2)$ 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 surface is the value function

$$V(T_1, T_2) = Y(T_1, T_2, B^*(T_1, T_2))$$

and it is analogous to blue ridge line in Panel B. This curved triangle is the upper envelope of all the production frontiers across all values of B, where individual surfaces of $Y(T_1, T_2, B_n)$ are exactly tangent to the value function for those positions in the climate space (T_1, T_2) where $B_n = B^*$, as was described by Equation 22. For example, the black translucent surface labeled $Y(T_1, T_2, B_2)$ lies below V for almost all positions in $\mathcal{C}$, but defines the maximum value obtainable for the small band labeled $B^* = B_2$, where the surfaces are exactly tangent. This tangency is the manifestation of Equation 22 (and the Envelope Theorem) which allows the shape of the value-function surface V to be measured, locally, without consideration for any adaptive adjustment of B. For an economy that is initially positioned at $V(T_1, T_2) = Y(T_1, T_2, B_0^*)$ and then is perturbed by idiosyncratic two-dimensional weather variation $(\tilde{T}_1(t), \tilde{T}_2(t))$, changes in observed output $Y(\tilde{T}_1, \tilde{T}_2, B_0^*)$ with respect to $\tilde{T}_1(t)$ and $\tilde{T}_2(t)$ will equal corresponding changes in the value function $V(T_1, T_2)$ with respect to T_1 and T_2. In panel F, the color of each position in the 2-simplex depicts the value function at that location in the climate space, analogous to the height of the curve in Panel C. Note that the height of the simplex (T_3) is not an independent dimension because it is fully determined by the first two dimensions (T_1, T_2).

Importance of nonlinearity for capturing adaptation empirically Both examples above demonstrate how idiosyncratic variations in weather around some baseline climate can be used to identify the local structure of the value function in the neighborhood of an initial climate $\mathbf{C}_0$, since in both cases adaptive adjustments have exactly no influence in the slope of a tangency plane. However, for non-marginal changes in climate, adaptive adjustments to $\mathbf{b}$ may be large, causing local derivatives of the value function $\frac{\partial V(\mathbf{C}_0)}{\partial \mathbf{c}}$ to change as the climate moves to positions in $\mathcal{C}$ far from $\mathbf{C}_0$. This is seen in both examples above, where the value surface V exhibits *curvature*, some of which is likely due in part to changing values of $\mathbf{b}^*$[8]. Thus, in order to ensure that any value function of climates identified by local variations in weather captures adaptations, it is important that a large number of marginal effects of weather (local tangencies) are estimated at different baseline climates (equivalent to evaluating Equation 25 at many different $\mathbf{C}$) prior to integration. This will allow the resulting value function to exhibit curvature. 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. This would have the effect of causing the one-dimensional value function in Panel C of Figure 2 to be linear, and the two dimensional value function in Panel E of Figure 2 to be a flat plane. These constant marginal effects models might reasonably approximate $V(\mathbf{C})$, although it is difficult to be certain whether this is the case *ex ante*. In our empirical implementation in the US, we explore when constant marginal effects are a good approximation and when nonlinearity is critical for adequately capturing adaptive responses to climate.

4 Empirical implementation for the modern United States

In our empirical analysis of the temperature climate for the US, we define $\psi(.)$ 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 extreme top and bottom bin 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 climate of a county i is then $\mathbf{C}_i$, a 16-element-long vector describing the expected count of days in each temperature bin less one. $\mathbf{C}_i$ is thus the position of that county on the 16-simplex $\mathcal{C}$.

Our objective is to empirically estimate the value of repositioning the climate $\mathbf{C}_i$ to some nearby

[8]It is theoretically possible for income to respond while $\mathbf{b}$ remains perfectly stationary for non-marginal changes in climate. This occurs if the PFF and better-than set evolve so that the separating hyperplane between the feasible and better-than sets rotates around a fixed $\mathbf{b}^*$. However, this seems like an unlikely scenario in most contexts.

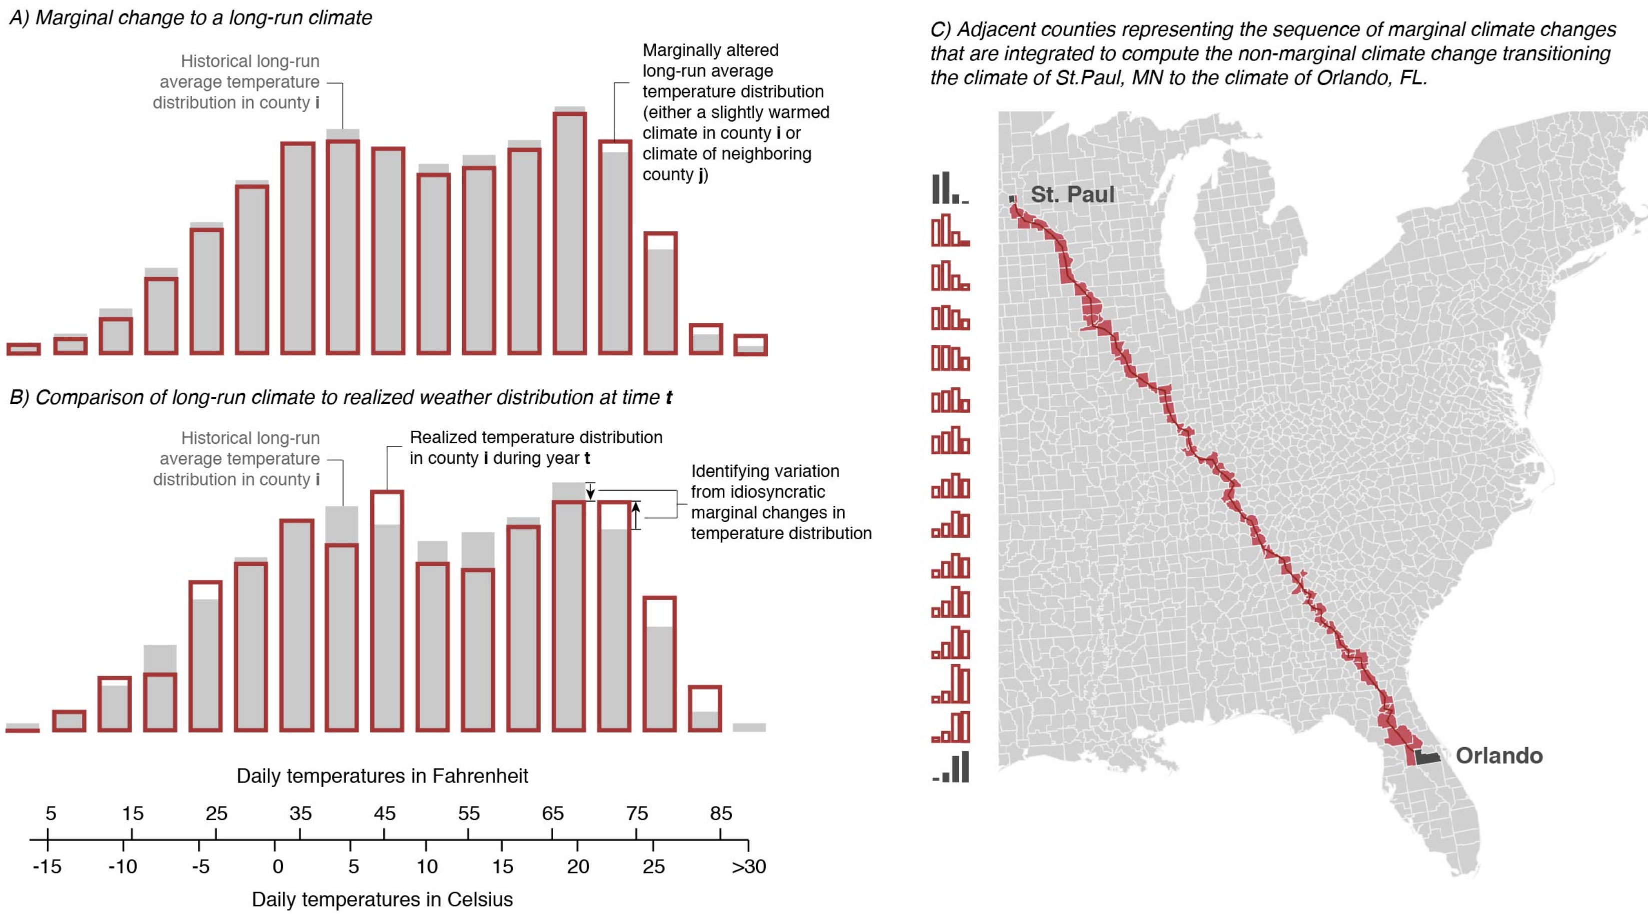

Figure 3: (A) Comparing the average temperature distribution of a county i (gray), summarized by a 17 bin histogram (bins are 3°C wide) that defines a position $\mathbf{C}_i$ in the 16-simplex climate space $\mathcal{C}$, with a slightly altered distribution (red outline) that could represent a permanent change to the climate at i or the climate of a neighboring county. (B) Comparing the climate $\mathbf{C}_i$ (gray) with a realized distribution of temperatures at i and time t summarized by the weather vector $\mathbf{c}_{it}$ that also lies in the climate space $\mathcal{C}$. Idiosyncratic deviations in $\mathbf{c}_{it} - \mathbf{C}_i$ identify the local structure of the value function $V(\mathbf{C})$ in the neighborhood of $\mathbf{C}_i$. (C) Example integration used to compute the difference in value of the climate in St. Paul, MN from that of Orlando, FL (both black). A sequence of changes in climate between neighboring counties are each marginal (cartoon temperature distributions on left) and trace out a line integral through the space of possible climates $\mathcal{C}$. Each marginal change in $V(\mathbf{C})$ is exactly identified by within-county variation in weather in the neighborhood of that position $\mathbf{C}$. Integration of these locally estimated marginal effects, traced along the red path through counties, allows calculation of the non-marginal difference between the value of climate in St Paul and Orlando (Equation 26).

location $\mathbf{C}_i + \mathbf{C}'$. To measure the value of such a permanent marginal distortion in the climate, our empirical strategy exploits the result in Equation 22 and measures the marginal value of temporary distortions to the weather vector $\mathbf{c}_{it}$ within location i over time, which is equal to the marginal value of the permanent distortion. 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 the 16-element disturbance vector ξ_{it} may have a complex structure and need not necessarily have mean zero.[9] As ξ_{it} varies randomly over time, the position of $\mathbf{c}_{it}$ will "explore" the neighborhood of $\mathcal{C}$ surrounding the position $\mathbf{C}_i$, allowing us to estimate the local marginal effect of these changes. Panel B of Figure 3 displays an example realization of $\mathbf{c}_{it}$ overlaid on the climate $\mathbf{C}_i$. 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 (Schlenker and Roberts, 2009; Deschênes and Greenstone, 2011).

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 26). If there is an adjacent county j such that $\mathbf{C}_j = \mathbf{C}_i + \mathbf{C}'$, then we can then exploit inter-temporal variation in $\mathbf{c}_{jt}$ to estimate the local structure of $V(\mathbf{C}_j)$ and thus compute $V(\mathbf{C}_j + \mathbf{C}')$ via another integration step. We can repeat this procedure, so long as adjacent counties are not "too far" away from one another in the space $\mathcal{C}$, and thereby compute the difference in value $V(\mathbf{C}_j) - V(\mathbf{C}_i)$ for arbitrary pairs of i and j, even if the change in climate is non-marginal.

Panel C 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 trace a path 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 integrate to follow the shape of the surface from $V(\mathbf{C}_i)$ to $V(\mathbf{C}_{i+1})$ through this space. It is worth noting that the exact path taken from St. Paul to Orlando should not matter, so long as the $V(\mathbf{C})$ is sufficiently smooth and counties are sufficiently "near" one another in $\mathcal{C}$, where "near" is relative to the curvature of $V(\mathbf{C})$. If $V(\mathbf{C})$ is relatively linear in $\mathbf{C}$, then extrapolation of $V(\mathbf{C})$ between counties distant in $\mathcal{C}$ may be reasonable.

4.1 Regression specification

We construct our empirical specification to closely reflect the theoretical structure of the climate value function such that the vector of regression coefficients we recover directly characterize the marginal product of climate in Equation 15. We combine Equations 15, 22, and 24, setting the benchmark climate to $\mathbf{C}_i$, and noting that the level of the value function at $\mathbf{C}_i$ is also conditional on

[9] 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, although there is no issue if this relationship holds.

a variety of other location and time specific factors unrelated to the climate $\mathbf{z}_{it}$, such as geography, history, and human capital endowments:

$$\begin{aligned} Y(\mathbf{C}_i + \mathbf{C}'_{it}, \mathbf{b}^*) = V(\mathbf{C}_i + \mathbf{C}'_{it}) &= V(\mathbf{C}_i|\mathbf{z}_{it}) + \int_{\mathbf{C}_i}^{\mathbf{C}_i+\mathbf{C}'_{it}} \frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}} \cdot \mathrm{d}\mathbf{C} \\ &\underset{\mathbf{C}'\to 0}{=} V(\mathbf{C}_i|\mathbf{z}_{it}) + \int_{\mathbf{C}_i}^{\mathbf{C}_i+\mathbf{C}'_{it}} \frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{c}} \cdot \mathrm{d}\mathbf{c} \\ &= V(\mathbf{C}_i|\mathbf{z}_{it}) + \int_{\mathbf{C}_i}^{\mathbf{C}_i+\mathbf{C}'_{it}} \left[\tfrac{\partial V}{\partial \mathbf{c}_{k=1}}, \quad \cdots, \quad \tfrac{\partial V}{\partial \mathbf{c}_{k=K}}\right]\Big|_{\mathbf{C}} \cdot \mathrm{d}\mathbf{c} \\ &\underset{\xi\to 0}{=} V(\mathbf{C}_i|\mathbf{z}_{it}) + \underbrace{\left[\tfrac{\partial V}{\partial \mathbf{c}_{k=1}}, \quad \cdots, \quad \tfrac{\partial V}{\partial \mathbf{c}_{k=K}}\right]\Big|_{\mathbf{C}_i}}_{\frac{\mathrm{d}V(\mathbf{C}_i)}{\mathrm{d}\mathbf{C}}} \cdot \xi_{it} \end{aligned} \tag{27}$$

where the last equality holds exactly for sufficiently small changes in ξ since it is the first order Taylor expansion of $V(\mathbf{C}_i + \xi)$. The first term is simply the total income that the economy at i would obtain at time t conditional on covariates $\mathbf{z}_{it}$ if the climate remained at $\mathbf{C}_{it}$—in the data we model this term parsimoniously using vectors of county fixed-effects, year fixed-effects, and an auto-regressive term. The second term in Equation 27 is the inner product between the K-dimensional gradient of the value function, evaluated at $\mathbf{C}_i$, and the disturbance vector ξ_{it} that describes how weather conditions at it deviate from characteristic weather conditions $\mathbf{c} = \mathbf{C}_i$. The gradient of the value function is the object of interest because it is the marginal product of the climate at $\mathbf{C}_i$, and it can be recovered from the data as coefficients to the terms in ξ_{it}. Importantly, because the marginal product of climate is identified locally, in the neighborhood of $\mathbf{C}_i$, it is possible that marginal products identified at different positions in $\mathcal{C}$ are not identical. This may occur if, as was shown in the illustrative examples above, populations adapt to climates in ways that alter its local marginal product. To account for this, 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 $\mathbf{C}$ in $\mathcal{C}$, thereby fully capturing all effects of adaptation. However, for clarity of exposition we begin by first presenting a linearized model where the marginal product of each K dimension of the climate is held fixed throughout $\mathcal{C}$. We then demonstrate how allowing for curvature in the value function alters these results.

We estimate the marginal product of climate by constructing an empirical analog to Equation 27, where the dimensions of climate include both the daily temperature and precipitation distributions for both the current and past year, although our main focus is on the effect of current temperatures since these other dimensions of the climate appear to have little effect on the value function. As mentioned above, we also account for within-county autocorrelation, unobserved heterogeneity across counties, and nonlinear time trends. Specifically, using our panel of US counties, we estimate

$$Y_{it} = \rho Y_{i,t-1} + \mu_i + \theta_t + \sum_{h=1}^{H} \left[\sum_m \left[\beta^{mh} (\tilde{T}^m_{it})^h + \gamma^{mh} (\tilde{T}^m_{i,t-1})^h \right] \right] + \sum_g \left[\zeta^n \tilde{P}^g_{it} + \eta^n \tilde{P}^g_{i,t-1} \right] + \epsilon_{it}, \tag{28}$$

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 geography. θ_t is a set of year fixed effects that flexibly account for common trends, such as technological innovations or trends in climate, and year-specific shocks, such as abrupt changes in energy prices. The model is allowed to be nonlinear in each dimension of the temperature climate up to the order H, to allow for curvature in the value function, which is crucial for fully accounting for adaptation[10]. In our linearized model, we set $H = 1$, which constrains the marginal effect of a hot or wet day to remain constant throughout the support of the weather data. We then re-estimate Equation 28 allowing for curvature in every dimension of $\mathcal{C}$ by setting $H = 3$. 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 in the mth temperature bin. Each interior temperature bin is 3°C wide. We define $\tilde{T}_{it}^{m=1}$ = the number of days when $T_d < -15$°C, $\tilde{T}_{it}^{m=2}$ = the number of days when $T_d \in [-15, -12)$°C, $\tilde{T}_{it}^{m=3}$ = the number of days when $T_d \in [-12, -9)$°C, and so on. The top ($m = 17$) bin counts days with $T_d \geq 30$°C=86°F. The $m = 11$ bin for $T_d \in [12, 15)$°C = $[53.6, 59)$°F is the omitted category.[11] The coefficients β^{mh} are the parameters of interest, as they characterize the marginal effect on Y of an additional day in the mth temperature bin, relative to a day with temperatures in the omitted category. In the linear model where $H = 1$, these coefficients are intuitive to interpret directly as they are the marginal effect of an additional day and the vector of coefficient estimates $[\hat{\beta}^{m=1}, ...\hat{\beta}^{m=16}]$ is directly interpretable as the gradient vector for the value function, the marginal product of climate $\widehat{\frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}}}$. When the order of the model is increased to $H = 3$ in order to fully account for adaptation, interpretation becomes more nuanced. The count of days in each mth temperature bin is accounted for in the model with the polynomial

$$\beta^{m1}\tilde{T}_{it}^{m} + \beta^{m2}(\tilde{T}_{it}^{m})^2 + \beta^{m3}(\tilde{T}_{it}^{m})^3,$$

which allows for the marginal effect of each additional day in the mth temperature bin to evolve like the derivative of this polynomial with respect to the count of days in the mth bin:

$$\beta^{m1} + 2\beta^{m2}\tilde{T}_{it}^{m} + 3\beta^{m3}(\tilde{T}_{it}^{m})^2.$$

For ease of interpretation, we display the full polynomial below, where the importance of non-constant marginal effects becomes clear. However, it is important to note that, unlike the case where $H = 1$, the vector of coefficients (which is now three times longer) cannot be directly

[10] We find that the role of precipitation is essentially zero, even in a linear model, so we do not present models that account for curvature in the precipitation subspace of $\mathcal{C}$.

[11] For display purposes, coefficients on the two coldest temperature bins ($T_d < -15$°C and $T_d \in [-12, -9)$°C) are not shown in figures but are included in tables. Generally, effects in these bins are highly uncertain and not statistically different from zero because there are few observations at these extremely cold temperatures.

interpreted as the gradient vector of the value function because the first-order Taylor expansion used in Equation 27 is no longer exact.

$\tilde{P}^g$ is defined similarly for daily precipitation across 12 bins. Each bin spans 40mm of daily precipitation, with the bottom bin corresponding to no precipitation and the top bin corresponding to precipitation $>$ 400mm in a day. Because temperatures and precipitation are, on average, serially correlated within a county, we include lagged values for all $\tilde{T}^m$ and $\tilde{P}^g$ variables to capture any possible direct effects that weather in the prior year might have on current output. For example, low rainfall in a prior year might reduce the quantity of groundwater available for irrigating crops in the current year.

The variable ϵ_{it} is a disturbance term that we assume may be arbitrarily correlated between counties within a state-by-year as well as within a given county over time. To account for this, we estimate standard errors that are clustered in two dimensions (Cameron, Gelbach, and Miller, 2011): within state-by-years and within counties. This approach accounts for both spatial correlation across contemporary counties within each state and autocorrelation within each county.[12]

Finally, $Y_{i,t-1}$ is a lagged dependent variable with serial correlation coefficient ρ. Including this term in the specification is important because there is substantial serial correlation in outcomes at the county level that is not accounted for by common trends. For example, the history of capital investments within a county affect production in subsequent years. It is known that one drawback of dynamic panel models, such as Equation 28, 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. $\leq$10 periods). We are not in this hazardous context, as our panel has 43 periods. We estimate that the magnitude of our potential bias is less than 5% of the magnitude of our point estimate, leaving us relatively unconcerned about this potential bias as it is far smaller than our uncertainty due to sampling error.[13] We opt to utilize OLS because it conveys many advantages, allowing us to account for spatial autocorrelation and avoid using weak or potentially invalid instruments. For completeness, we show estimates without any lagged dependent variable and continue to obtain our main result.

The average marginal effect of daily temperature in the mth bin (β^{mh}) is identified by Equation 28 if the exact number of days in that bin, relative to other years in the same county, are orthogonal to other potential confounders, conditional on all control variables. For example, the estimated effect of a 16°C day is identified by comparing a county to itself across years when the number of 16°C degree days was slightly different. Weather has systematic patterns in each location that are absorbed by county fixed effects. Random variations in those patterns give rise to small distortions in the distribution of daily temperature across years that we exploit for inference. Our estimates of each β^{mh} are identified off of these random disturbances at each point in the temperature

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

[13]Nickell (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.

distribution within a location.[14] We follow Deschênes and Greenstone (2007) and Schlenker and Roberts (2009) in assuming that these detrended year-to-year random variations within each county are uncorrelated with year-to-year variations in other important factors that affect income but are not caused by weather.

4.2 Data

We match weather and income data at the county level for the lower 48 states during the period 1969-2011. All income measures are inflation-adjusted to 2011 dollars and converted to per capita terms. Summary statistics for key variables are presented in Table A1.

Weather data To measure daily maximum and minimum temperatures as well as precipitation, we use daily surface data from the National Climatic Data Center (NCDC).[15] We match weather stations to counties using each station's reported latitude and longitude. We omit observations where the maximum or minimum temperature exceeds 60 degrees Celsius or is lower than -80 degrees Celsius, as these are likely errors. If there are multiple stations within a county, we average their measures for each day. Our preferred measure of daily temperature is a simple average between the maximum and minimum temperatures, which is the standard measure for average temperature during a 24-hour period.[16] As discussed in Auffhammer et al. (2013), weather station data is often incomplete, sometimes due to mechanical failures, political events, or financial constraints. We drop county-by-year observations that do not have a complete set of daily weather observations. This results in a reduced sample size, with coverage that is displayed in Appendix Figure A1. Thus, our results represent the average effect of temperature on income, conditional on whatever circumstances allow counties to provide a complete record of daily weather within a single year.

Income data To measure income, we use Regional Economic Information System (REIS) data, published by the Bureau of Economic Analysis (BEA). The BEA, in turn, uses a variety of sources to construct these measures.[17] The most inclusive income measure at the county level is *total personal income*. It encompasses all sources, including all types of labor income; proprietors' income; dividends, interest, and rent payments; and government transfer payments. A subset of personal income, *earnings*, includes only wages and salaries, other labor income, and proprietors' income. In turn, wages and salaries include tips, commissions, bonuses, and any "pay-in-kind" provided by an employer. They are measured before any deductions are taken and are derived from

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

[15] Publicly available from `ftp://ftp.ncdc.noaa.gov/pub/data/ghcn/daily/by_year/`.

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

[17] For further details, see `http://www.bea.gov/regional/pdf/lapi2010.pdf`.

reports filed by employers to comply with unemployment insurance (UI) laws.[18] Total personal income is reported on a place-of-residence basis, while wage and salary payments and other income components are reported by place of work. The residence adjustment is made using US Census estimates of worker commuting behavior. As a result, the components of personal income can sometimes exceed total personal income.

Measures of *farm income* in REIS are derived from United States Department of Agriculture (USDA) estimates, which are based on sample surveys, Agricultural Census data, and administrative data.[19] A distinction is made between gross farm income, which includes inventory sales, and net farm income, which does not. However, additions to inventories are included in the net farm income measure. Importantly, the *net farm income* measure we use also includes transfers such as subsidies, crop insurance, and disaster payments. Our measure of gross farm income is *cash receipts from marketing crops*.

5 The marginal product of temperature

5.1 Results for an affine value function with constant marginal products

We first present results that assume the value function $V(\mathbf{C})$ is a flat hyperplane spanning the 16-simplex $\mathcal{C}$ defining all possible contemporaneous temperature distributions, i.e. $H = 1$ in Equation 28. This model provides a good first-order approximation of the marginal product of climate for most county-years in our sample.

Main result Panel A of Figure 4 presents the first order approximation for the marginal product of daily temperature with respect to total annual income per capita in US counties. Specifically, the figure displays the vector of ordered coefficient estimates for contemporaneous daily temperatures that is equal to the gradient vector of the value function $V(\mathbf{C})$

$$[\hat{\beta}^{m=1}, ...\hat{\beta}^{m=16}] = \frac{\widehat{\mathrm{d}V(\mathbf{C})}}{\mathrm{d}\mathbf{C}}$$

which is precisely the marginal product of climate. As noted in many previous studies[20], the display of this high-dimensional vector in this way is convenient because it can be interpreted as a "dose-response" function where values indicate the marginal effect on annual income of exchanging a single day from the omitted category with a day of the indicated temperature. However, this

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

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

[20] We believe Deschênes and Greenstone (2011) were the first to point this out.

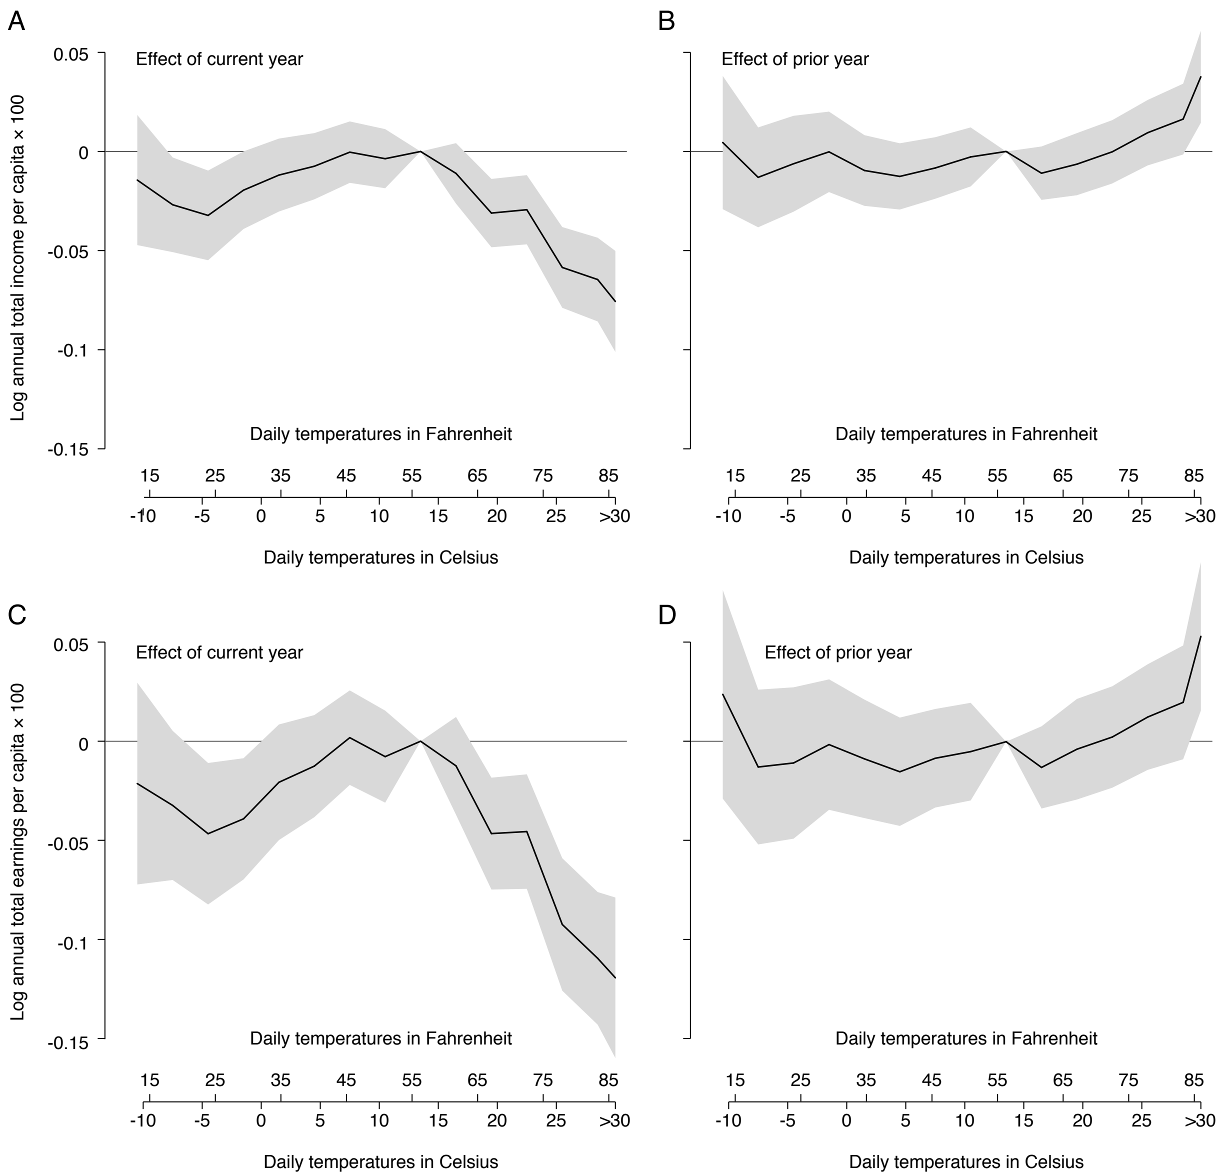

Figure 4: (A) The effect of distorting the temperature distribution by shifting a single daily average temperature on log annual total personal income per capita×100 (i.e. percentage points) in US counties for 1969-2011. For reference, an average day contributes $\frac{1}{365} = 0.27\%$ of annual income. Shaded area denotes 95% confidence intervals. (B) The effect of daily average temperatures in the prior year on income per capita. (C) Same as Panel A, but for total earnings per capita. (D) Same as Panel B, but for total earnings per capita. Panels A and B are estimated simultaneously in a single regression model. The same is true for Panels C and D.

convenient interpretation is no longer possible in the higher-order model ($H = 3$) that fully accounts for adaptation, presented in the next section.

We find that 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.[21]

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

Temporal displacement Panel B of Figure 4 displays the estimated effect of daily temperatures on annual income per capita the *following* year. We estimate these effects jointly with the contemporaneous effect shown in panel A, specified as γ^m in Equation 28. Except for the single coefficient in the hottest temperature bin (> 30°C), we do not observe any statistically significant effect of daily temperatures on income the following year. It is possible that the significant coefficient in the top temperature bin is spurious; because we are testing sixteen coefficients, it would not be unlikely for one to be spuriously significant. However, it is also possible that this effect is meaningful and indicates that some of the income lost from the hottest days is displaced into the following year. The estimated magnitude of this positive lagged effect is half the magnitude of the negative contemporaneous effects, indicating that roughly half of the income loss from the hottest days might be made up in the following year.[22] There is no statistically significant evidence for temporal displacement of income for days below 30°C.

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

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

[22] In 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 the effect daily temperatures on income, earnings, and farm income. All effects are scaled to the estimated effect of treatment in lag year zero (100%). Black lines are the effect of any disturbance to county income (temperature or otherwise) in the AR1 benchmark model for the decade following the shock—identification of this autoregressive structure is almost entirely dependent on non-temperature idiosyncratic income shocks. Colored lines are the directly estimated lagged effects of temperature effects in the top four bins that occur in the initial year, using a model that omits the AR1 structure. The inter-temporal structure of temperature impulse-responses is extremely similar to the AR1 characteristic structure exhibited in response to non-temperature impulses.

Intertemporal dynamics A key question in studies of climate's influence on production is whether temperature fluctuations affect income levels or growth rates, as the latter accumulate and can cause much larger income reductions in the long run. For example, early work by Nordhaus and Yang (1996) assumed that higher temperatures affected the instantaneous levels of production only, while country-scale empirical analysis by Hsiang (2010), Dell, Jones, and Olken (2012) and Burke, Hsiang, and Miguel (2015) find evidence that growth rates of national incomes are influenced by temperature changes, generating level effects that persist after temperatures return to earlier levels. For an economy in general equilibrium that adjusts to changes in climate (recall Figure 1), temporary changes in temperature will have some persistence if the allocations $\mathbf{b}^*$ in one period affect the endowment and allocations in subsequent periods. This could occur, 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 transmitting information about historical climate conditions into future periods (Burke, Hsiang, and Miguel, 2015). Furthermore, accounting for these intertemporal effects is important when trying to fully account for adaptation in the presence of credit markets, as some of the costs of reallocating $\mathbf{b}^*$ might be deferred to future periods.

In our setting, where we consider a large number of the "small" macro-economies of counties—which are known to have different dynamic properties than national economies—we observe that annual income in US counties is serially correlated, so a disturbance to income in county i in year t will have an indirect effect on income for i in year $t+1$. For average personal income in our lagged-dependent variable model (LDV, Equation 28), we estimate that $\hat{\rho} = 0.825$, implying that county income (unlike national income) does not have a unit-root, but instead that $Y_{i,t+1} = 0.825Y_{it} + X_{it}\beta + \epsilon_{it}$. Thus, an income loss of \$1 in year t will result in an income loss of \$0.825 in year $t+1$, \$0.68 in year $t+2$ and so on, relative to a counterfactual income trajectory where no loss was suffered in year t. Importantly, however, this pattern of auto-correlation is primarily identified off of variations in the idiosyncratic disturbance term ϵ for prior periods—not from changes in county temperature in prior periods—as most of the income variation is not driven by temperature. Therefore, it is important to check whether *temperature-induced* changes in income exhibit similar dynamics. It is possible that endogenous adjustments to state-variables made in response to temperature are more or less persistent than adjustments made in response to other idiosyncratic income changes.[23] To investigate this, we estimate a variant of Equation 28 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 similar to other types of income disturbances, than the structure of these lags should be similar to the negative-exponential structure of persistence indicated by $\hat{\rho}$, when the latter is no longer captured by the model. To implement this comparison, we examine the contemporaneous and lagged coefficients on key temperature bins to the analogous effect on future periods that we be captured by successive multiplication of the autoregressive coefficient.

The results of this exercise for log income per capita, log earnings per capita, and farm income

[23] We thank James Stock for this suggestion.

are shown in Figure 5. To ensure comparability, all contemporaneous effects of temperature are normalized to 100 and subsequent lags can be interpreted as a percent of the contemporaneous effect. The thick black line corresponds to the repeated application of the estimated autoregressive coefficient to the effects of temperature in our benchmark LDV model, and is smooth by construction. The blue, yellow, orange, and red lines correspond to the $21-24°C$, $24-27°C$, $27-30°C$, and $> 30°C$ temperature bins, respectively. Because 252 additional parameters are needed to characterize persistence (9 additional lags for 16 temperature bins and 12 precipitation bins), the unrestricted lagged coefficients on temperature bins are much noisier than the estimates from the LDV model. Nonetheless, they clearly decline over time with an average structure that resembles the negative exponential decay. Temperature-driven income responses appear to be persistent. For total income and earnings, the delayed effects of the hottest three bins appear to decay more rapidly than in the LDV model, although this difference is mostly explained by the small delayed positive response from hot lagged temperatures in the LDV model (recall Figure 4B and D). For farm income, the initial persistence of temperature-driven income changes tend to be substantially lower than in the LDV model, but the ratios for lags larger than three years are very similar to the autoregressive term (and close to zero). It 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, perhaps similar to the original intuition of Nordhaus and Yang (1996), although some of this recovery is likely due to crop insurance indemnities paid out by the federal government following hot temperatures (see the Online Appendix).

The finding that high-temperature-driven income losses exhibit persistence, although perhaps less so than other idiosyncratic disturbances, has two important implications. First, changes in the current temperature distribution of a county change the income trajectory of that county in subsequent periods, even if the temperature distribution returns to initial values in subsequent periods, indicating that some of the adaptive adjustments of $\mathbf{b}^*$ are in durable state-variables that themselves alter the evolution of the economy. This result thus provides county-level support for prior macro-economic findings that climatic conditions alter growth grates. Second, evidence of persistence indicates that the total NPV cost of a small change in the temperature distribution is larger than the contemporaneous marginal product would suggest. Using a 3% discount rate to integrate the transient evolution of the economy, a change in total income due a shift in the temperature distribution produces an NPV of income changes that is 5.0 times as large as the magnitudes displayed in Figure 4A.[24] Below, we do not present this NPV and instead focus on the

[24]The discounted sum of these income losses between the time of the temperature event and ∞ represents the net present value (NPV) of lost income attributable to the temperature event. By computing the NPV of the difference in county i's income (ΔY_{is}^m) at each moment s that was attributable to the temperature event $\tilde{T}_{it}^m = 1$ at time t, we can estimate the full net present value of a day with $\tilde{T}_{it}^m = 1$. Using a discount factor δ, this is

$$NPV(total\_lost\_income) = \sum_{s=t}^{\infty} \delta^{(s-t)} \Delta Y_{is}^m \approx \sum_{s=t}^{\infty} \delta^{(s-t)} \hat{\rho}^{(s-t)} \hat{\beta}^m = \frac{1}{1-\hat{\rho}\delta} \hat{\beta}^m.$$

Thus, the NPV of the altered income trajectory is a linear scaling of coefficients by $\frac{1}{1-\hat{\rho}\delta}$. Using a discount factor $\delta = 0.97$ (implying an annual discount rate of 3%) and $\hat{\rho} = 0.825$, we estimate this scaling factor to be 5.01. In NPV

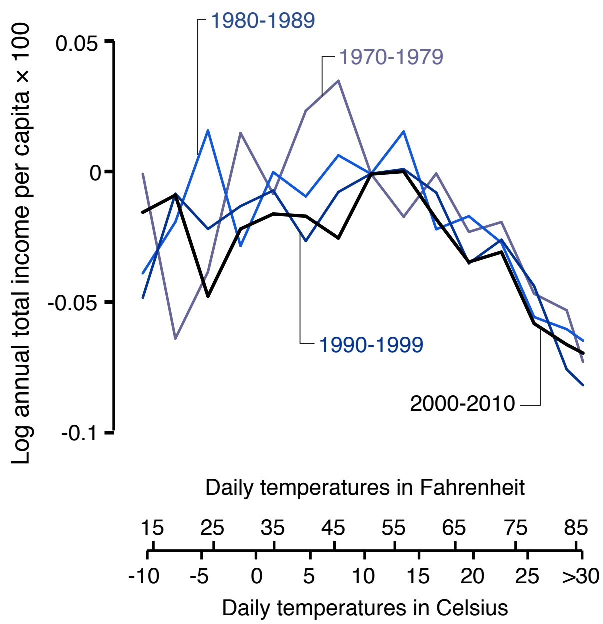

Figure 6: The stationary marginal product of climate over time. The gradient of the value function for the temperature distribution (as in 4A) using decade-long subsamples of the data.

contemporaneous marginal product of climate, although we do account for this transient response in 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. Some economic consequences of climatic conditions have surely changed over the period 1970-2010. 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., 2013). Meanwhile, other technology, such as the heat-tolerance of maize, has remained essentially stationary over the same period (Roberts and Schlenker, 2011). Importantly, prior analyses that document evolving sensitivity of specific impacts over time do not account for the total net cost of those adaptations which, presumably, are rising over time as populations increasingly exploit new technologies. Because the marginal costs and gains from additional adaptation must equalize at $\mathbf{b}^*$ at any moment in time, these previously unaccounted for costs must be similar in magnitude (on the margin) to gains from adaptation. Because our estimate of the gradient in the value function captures all market costs and benefits[25], we capture these adaptation costs, which, as mentioned before, are simply opportunity costs due to re-optimization

terms, this implies that a warm day at 22.2°C (72°F) costs $5.01 \times -0.000294 = -0.00147$ log points of annual per capita income (relative to a 13°C (55.4°F) day) because, in addition to altering contemporaneous income, it alters a county's future income trajectory. A hot day exceeding 30°C (86°F) is estimated to cost $5.01 \times -0.000757 = -0.00379$ log points of annual per capita income in NPV terms. Recalling that a randomly selected day is responsible for $\frac{1}{365} = 0.00274$ log points of annual income, the NPV of the total cost of a warm or hot day is roughly 0.5 and 1.4 days' worth of average income, respectively.

[25]We note that the analysis by Barreca et al. (2013) focuses on human mortality, the value of which is unlikely to be fully captured by the market.

of $\mathbf{b}^*$. In addition, many aspects of the US economy outside of temperature-related technologies have evolved during 1969-2011. Thus, taking into consideration both the previously unaccounted-for cost of adaptation as well as secular changes in the economy, it is not clear *ex ante* how the marginal product of climate should evolve over time.

To examine how the marginal product of climate has evolved, we re-estimate the response of income to temperature for each decade separately, shown in Figure 6. These estimates are noisier because each relies on a much smaller sample, but they do not differ substantively from our pooled estimate or from one another. In particular, the marginal effect of warm and hot days is essentially constant over time, indicating a remarkable stability of the gradient of the value function $V(\mathbf{C})$ in the $\mathcal{C}$-subspace governing those temperatures. This indicates that despite the possibly numerous adaptations to temperature that have occurred during the last half-century in the US—adaptations which may have substantially increased welfare—they did not fundamentally alter the marginal product of climate. High temperature days remain costly even during the first decade of the twenty-first century.

Seasonal consistency Our estimates reflect the average marginal product of climate given the current distribution of daily temperatures within the year. We examine whether these marginal products vary substantially across seasons by estimating a version of Equation 28 where separate contributions to annual income for temperature distributions in each quarter of the year are estimated simultaneously, by interacting temperature bins with dummy variables for each quarter (Jan-Mar, Apr-Jun, Jul-Sep, Oct-Dec). The results are shown in Figure 7A-D. 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, which is sensible because the bulk of variation in those dimensions of $\mathcal{C}$ occur during those quarters. A similar pattern hods 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 events are uncommon. Overall, these results suggest that the marginal product of climate, for those regions of $\mathcal{C}$ in which sufficient data allows estimation, is quite consistent across seasons in our setting.

Regional heterogeneity The marginal product of climate is the *local* gradient of the value function $V(\mathbf{C})$, measured in the vicinity of some position $\mathbf{C} \in \mathcal{C}$. Thus, if there is any curvature in value function (recall Figure 2), the local gradient vector—which must point in the direction of steepest ascent at all points along the surface $V(\mathbf{C})$—will *rotate* as $\mathcal{C}$ is traversed. Stated another way, the local slope of the value function will change based on the average climate of a location. Such curvature could result from endogenous adaptations to different climates, if populations select different $\mathbf{b}^*(\mathbf{C})$ in different climates (e.g. investing in air conditioners). But it could also result from features of the income-generating process that are beyond the control of agents, as the PPF may itself respond nonlinearly to temperature. For whatever reason, if there is such curvature, then the marginal product of climate will be different for different regions of the country, because

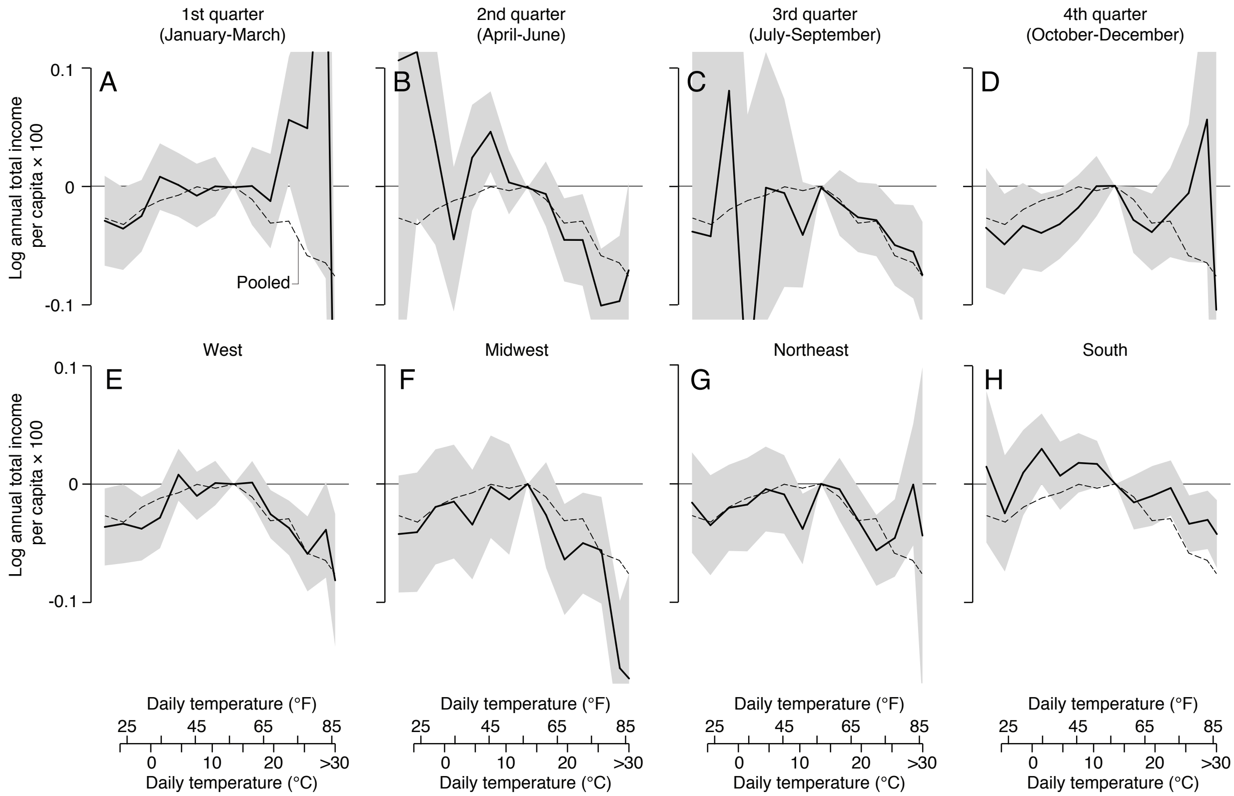

Figure 7: The marginal effect of climate estimated for sub-periods within each year and sub-regions of the country, compared to the pooled estimate (dashed, identical to Figure 4A). (A)-(D) Effects of temperatures experienced in each quarter of the year (estimated simultaneously), where annual income is allowed to respond to changes in the temperature distribution for all four quarters separately. Large confidence intervals generally indicate regions of the temperature distribution containing very few observations.(E)-(H) Effects of annual temperature distributions by region (estimated simultaneously).

counties in different regions will, on average, be located in different regions of $\mathcal{C}$. We provide *prima facie* evidence of curvature in $V(\mathbf{C})$ by examining the marginal product of climate for four major regions of the country, the West, Midwest, South, and Northeast (as defined by the US Census).

In Figure 7 we examine regional heterogeneity in the data and find that the overall structure of the response to daily temperature distributions is generally similar across the country, suggesting that curvature in the value function due to adaptation or the PPF 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 influence of high temperatures is not identical everywhere. In particular, high temperatures are most costly in the Midwest and least costly in the South. The high marginal cost of hot temperatures in the Midwest is likely due to changes in the PPF associated with the climate in the neighborhood of $\mathcal{B}$ where Midwestern economies tend to be located—with many resources allocated to agricultural production. In the South, low marginal costs of hot temperatures

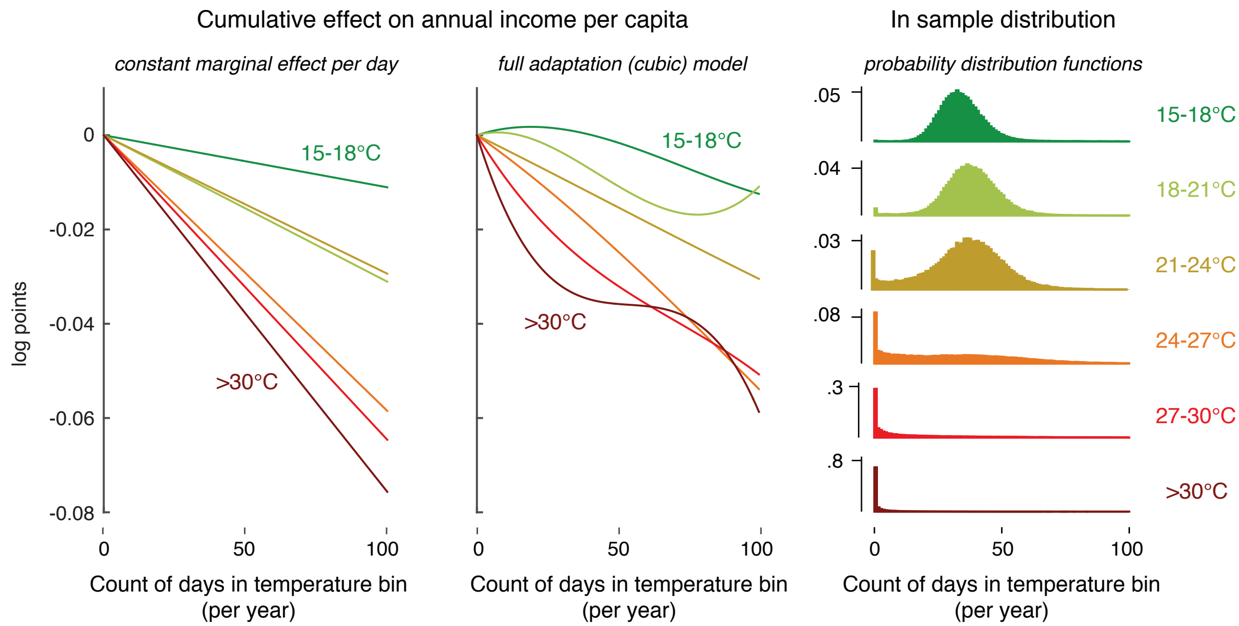

Figure 8: Comparison in the estimated effect of hot temperature days using an approach that assumes constant marginal effects of additional days (left) and an approach that accounts for adaptation (i.e. curvature in the 16-simplex $V(\mathbf{C})$) by allowing income to be a cubic function of the count of days in each temperature bin (middle). Histograms displaying the distribution of county-years across the support of each bin-specific response. Fully allowing for adaptation primarily allows for some concavity in the response to moderate temperature days (15-21°C) and convexity in the response to very hot days (>30°C). Responses of warm temperatures are well approximated by the constant marginal effects model.

appear related to adaptive allocations $\mathbf{b}^*$ due to the hot climate of the region. Later, we investigate possible reasons for these differences. But before investigating mechanisms directly, we illustrate the importance of curvature in the value function to net out effects of adaptive adjustments.

5.2 Results for a value function with curvature

We next present results that do not assume the value function $V(\mathbf{C})$ is flat. Rather, we allow the envelope of equilibrium allocations $\mathbf{b}^*$ on the PPF to trace out a curved surface spanning the 16-simplex $\mathcal{C}$ (analogous to Figure 2, but in 17 dimensions rather than three). In contrast to assuming the value function is linear, this approach allows the marginal effect of distorting to the temperature distribution by $\mathbf{C}'$ to have effects that depend on the initial position of a county's economy $\mathbf{C}_0$ within $\mathcal{C}$. Such an adjustment becomes important if, for example, the marginal effect of additional hot days declines as a population experiences higher numbers of hot days, perhaps because some adaptive allocations (such as air-conditioning) become increasing beneficial when more hot days are experienced. If this were the case, and hot days were still costly on average, then the value function would become convex with a less extreme gradient as economies moved toward the corner of the simplex $\mathcal{C}$ where hot days dominate the climate. We implement this approach by setting $H = 3$ in Equation 28, such that income becomes cubic in the count of days in each bin. This allows the shape of the value function to be quite flexible, since it may curve differently in each dimension of the temperature distribution. Because this approach accounts for all changes in marginal effects due to adaptation, we denote it the "full adaptation model," in contrast to the "constant marginal effect" model.

Figure 8 presents the results of the full adaptation model, contrasted with the constant marginal effects model, for all temperature bins above 15°C . Because the gradient vector of the value function is no longer constant, it cannot be plotted as a "response function," as in earlier sections. Instead, we display the total contribution of days within each temperature bin to annual income. The derivatives of these functions are the marginal effect of an additional day at that temperature, relative to the omitted 12-15°C bin. The right-most panel displays probability density functions over the count of days in each temperature bin within our sample, with colors and labels corresponding to the functions in the center and left panels.

In the constant marginal effects model (left-most panel of Figure 8), we see that annual income declines linearly in the count of days within each warm or hot temperature bin, with slopes that reflect the constant marginal effects in Figure 4A. In the full adaptation model (middle panel of Figure 8) we see that that marginal effects of modest but 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, with marginal effects that increase with the count 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 and identical to the constant marginal effects model. The concavity of the 15-18°C bin, with its peak between 22 and 23 days per year, indicates that most counties actually suffer very small losses to these temperatures, as the distribution of daily counts are clustered and centered only slightly right of this peak (right panel).

For the hot 27-30°C bin, the convexity is modest, with the first day in this bin reducing annual incomes -0.096 log points, the tenth day reducing income by -0.081 log points, and the thirtieth day reducing income by -0.054 log points. Since almost all county-years experience fewer than ten days in this category, this response remains generally 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 very hot day in the annual distribution of daily temperatures lowers annual incomes by -0.181 log points (recall that a single day contributes roughly 0.274 log points on average) while the tenth hot day a year lowers incomes only -0.125 log points and the thirtieth lowers incomes -0.039 log points. The value function is estimated to have an inflection point near 50 days per year above 30°C , although estimates in that region are extremely noisy since there is essentially no probability mass in that region: over 70% of county-years in our sample have only one or zero days in this bin, with only very hot and arid regions of the country, such as the deserts, exhibiting large number of days with such high average temperatures. Nonetheless, the flattening out of the value function for high counts of very hot days clearly indicates that populations with many hot days make large adaptive allocations of $\mathbf{b}^*$ such that hot days have limited marginal impact, although these reallocations appear sufficiently substantive that they have large average impacts on income.

To our knowledge, this estimate of the non-affine 17-dimensional value function represents the first characterization of the market value generated by the temperature distribution, accounting for all benefits and costs of adaptive adjustments captured by the market. Importantly, however, this result reveals 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 other properties of the data in our sample. As will be seen below, focusing on results that assume constant marginal effects substantially reduces the complexity of displaying and understanding properties of the value function, a simplification that becomes important as we explore additional heterogeneity in the data and try to understand the mechanisms underlying this result.

Importantly, 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. If marginal effects are assumed to be constant, than our model estimates will poorly approximate the true value function and we will mis-estimate the total cost of warming a climate. As we demonstrate below, this issue becomes especially important when we project income changes into a future climate change scenario using our estimated value function. Thus, when we consider the overall value of current and future climates, we account for curvature along every dimension of the value function surface.

Erroneous "folk theorem" on sign of forecast bias We make one final practical point that is revealed by our estimate for the fully nonlinear value function which we think is critical and applies widely throughout this empirical literature. Our nonlinear results in Figure 8 clearly indicate that populations adapt to high temperatures and that assuming constant marginal effects for these high temperatures will generate incorrect predictions about the effects of warming a population's climate. But an incorrect "folk theorem" salient in the climate literature states that projections using empirical estimates that assume constant marginal effects of temperature will necessarily generate an "upper bound" for the damage from warming, because adaptations will cause the actual marginal damages from warming to become nearer zero. We carefully demonstrate that this reasoning is incomplete below. The intuition for why is seen most clearly in the panels of Figure 8. Focusing on the response to very hot ($> 30°$C) days for clarity, we see that the constant marginal effects model does over-estimate the marginal damage from additional warming for highly adapted populations, consistent with the "folk theorem." However, this pooled model also *under-estimates* the marginal damage of additional warming for cooler and poorly adapted populations that presently have only zero or one day in the hottest temperature bin—i.e. the slope of the fully adapted response is actually *steeper* at zero days 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 bulk of counties—are assigned marginal effects that are too small in magnitude because they have been averaged with the marginal effects from a smaller number of hot and highly adapted locations. Thus, as counties with only zero or one very hot day (right panel) warm up, they initially descend down the fully adapted curve much more rapidly than they descend down the constant marginal effects line. Because in the US most economic activity at present 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, a fact that we demonstrate below.

6 Mechanisms

We attempt to understand the mechanisms that underly the nonlinear and multidimensional response of total income to the distribution of temperatures. Much research on the economic impact of temperature in the US has focused on farming, since the negative impact of adverse weather on crop yields is acute (Mendelsohn, Nordhaus, and Shaw, 1994; Schlenker, Hanemann, and Fisher, 2005; Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009; Welch et al., 2010; Fisher et al., 2012; Burke and Emerick, 2016). However, recent work has indicated that non-farm income sources outside of the US may also be affected by high temperatures (Jones and Olken, 2010; Hsiang, 2010; Dell, Jones, and Olken, 2012; Burke, Hsiang, and Miguel, 2015). These authors suggest that this non-farm effect may be driven by the well-documented productivity decreases of workers who are exposed to thermal stress (Mackworth, 1946; Froom et al., 1993; Seppanen, Fisk, and Lei, 2006). Furthermore, changes in aggregate demand could theoretically respond to temperatures, shifting equilibrium production and altering prices.

In the US, high daily temperatures are known to reduce yields of major crops (Schlenker and Roberts, 2009) as well as to reduce labor supplied among workers exposed to outdoor temperatures, which includes manufacturing (Graff Zivin and Neidell, 2014). These studies demonstrate that the productivity of crops and the quantity of labor supplied depend on daily temperature and suggest mechanisms that might explain our main finding. In our general equilibrium framework, these changes can be thought of as altering the structure of the PPF. However, without observing price changes, these studies alone are not conclusive. If local or regional prices change with temperature, then changes in production might not translate into changes in revenue. To consider whether these mechanisms might be contributing to the effect we document above, we directly examine how the agricultural and non-agricultural components of income respond to temperature. We then compare both the structure and magnitude of these responses to earlier results by Schlenker and Roberts (2009) and Graff Zivin and Neidell (2014). To facilitate comparison, we reproduce the main results of both studies in Figures 9A and 9D, 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. Lab studies indicate that the labor productivity response to temperature is qualitatively similar in structure to the response reported in Graff Zivin and Neidell (2014) (Mackworth, 1946; Froom et al., 1993; Seppanen, Fisk, and Lei, 2006). Thus, the total labor effects on income may be larger than the estimates in Graff Zivin and Neidell (2014) suggest, but the overall structure of the response should be similar.

6.1 Agriculture

To examine how crop losses contribute to our main result, we repeat our analysis, replacing the dependent variable with *log revenue from crop sales per capita*. In panel B of Figure 9, we plot the effect of hot days on annual income from crops and observe steep declines when daily average temperatures rise above 27°C. This structure is very similar to the yield response obtained by Schlenker and Roberts (2009) (Figure 9A). The slightly higher breakpoint of 29-32°C in that study and its steeper decline is likely because Schlenker and Roberts (2009) use hourly temperature, whereas our analysis uses daily averages. Because days with 24-hour average temperatures of 27°C are likely to have some hours above 29°C, we would expect to observe declines on days with average temperatures of 27°C in our analysis, even if crop yields do not deteriorate until the hourly temperature reaches 29°C. Thus, we interpret our results in Figure 9B as consistent with the crop yield response in Figure 9A reported by Schlenker and Roberts (2009). Our results suggest that higher crop prices do not dramatically offset yield losses caused by high temperature days; thus, reductions in yields translate into reductions in income.

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

In panel C of Figure 9, we examine how *net farm income per capita* (in levels) responds to daily temperature and find that it declines by $21.07 for each day above 30°C .[26] The structure of the response of total farm income differs somewhat from that of crop income: we observe lower farm incomes starting at temperatures around 20°C . We lack the data to determine precisely

[26] Net farm income is not amendable to a log model because many observations in the sample are negative. So we estimate an analogous model but with the outcome specified in levels.

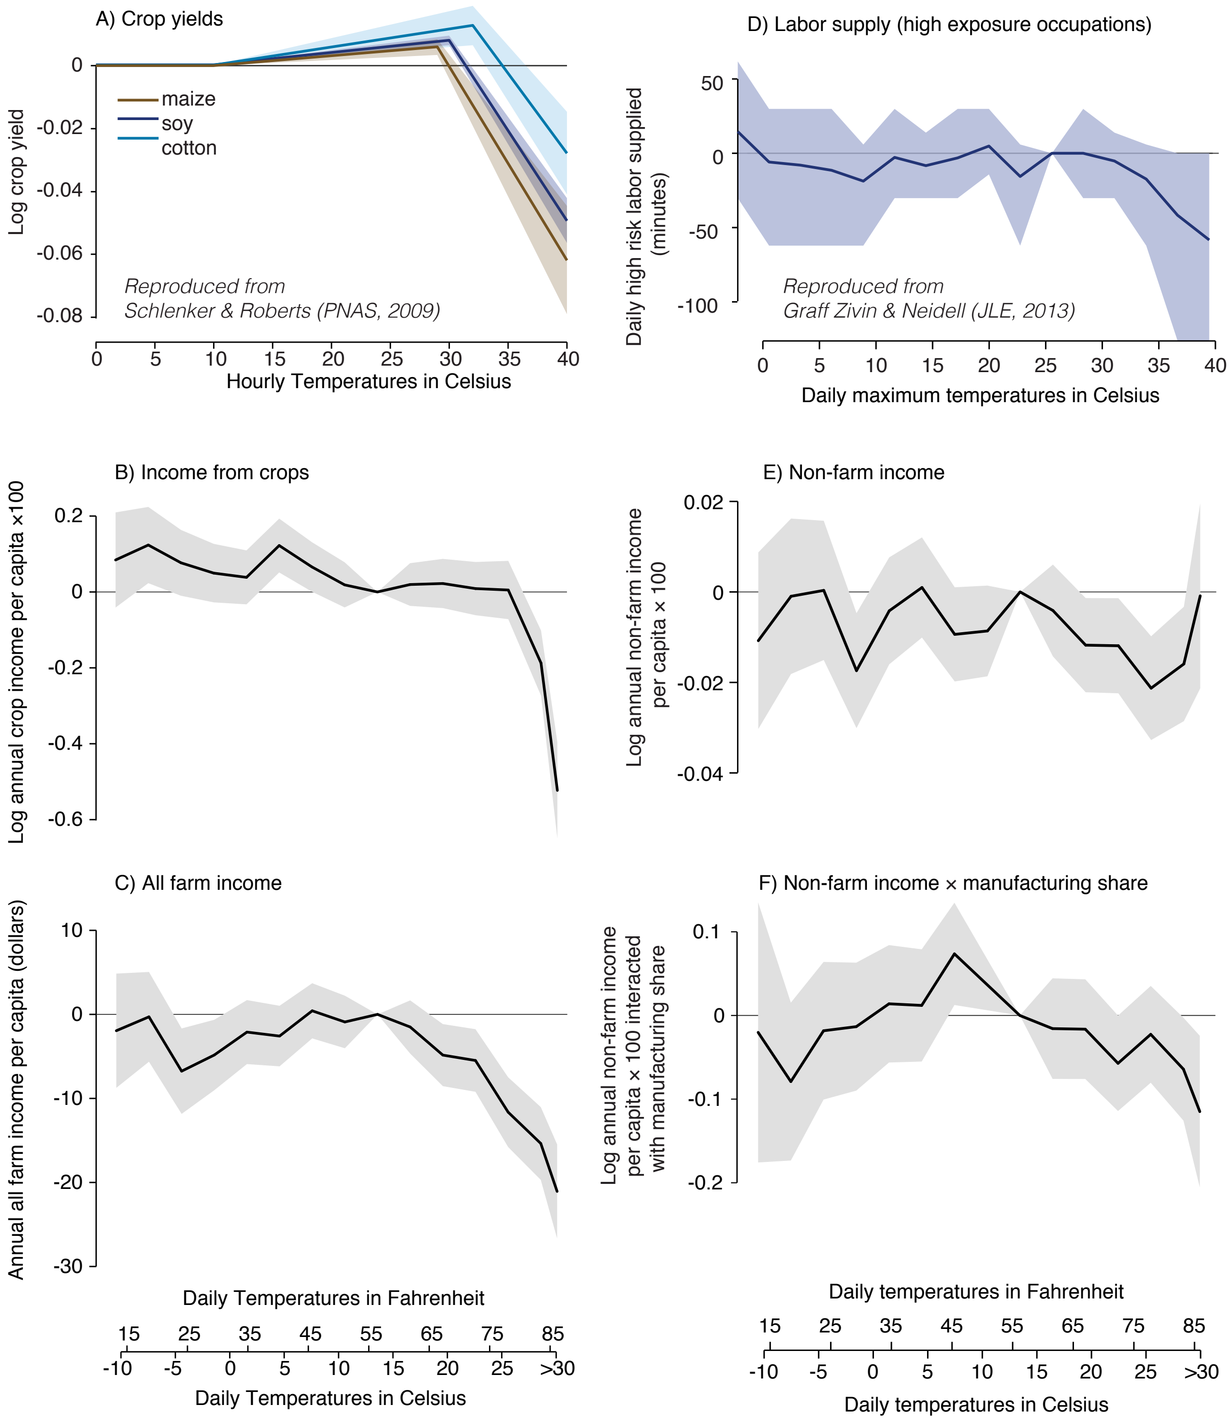

Figure 9: (A) Log annual crop yields vs. temperature during growing season for three major crops, reproduced from Schlenker and Roberts (2009). Yield effects are depicted as the effect of 24 hours at exact temperatures. (B) The effect of daily average temperature on log income from crops per capita, from this study. (C) The effect of daily average temperature on farm income per capita (in levels), from this study. (D) 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). (E) The effect of daily average temperature on non-farm income per capita, from this study. (F) The effect of daily average temperature on non-farm income per capita interacted with county-level measures of manufacturing income share (U.S. Census Bureau, 1969-2011).

the mechanism that could mediate effects at these lower temperatures. However, the observed structure is broadly consistent with the labor productivity response discussed below. Alternatively, farmers may be increasing expenditure on inputs to combat the negative impacts of temperatures on yields. Finally, it is also possible that other productive factors respond negatively to these lower temperatures.[27]

6.2 Non-farm income and manufacturing

We examine whether non-agricultural income might be playing a role by repeating our analysis on *log non-farm income per capita*, shown in Figure 9E. We see that non-farm income is relatively flat (albeit noisy) at low temperatures and then begins to decline systematically at temperatures above 15°C, the same breakpoint observed for total income (Figure 4). However, the magnitude of the effect on non-farm income is smaller, with temperatures at 25°C lowering annual non-farm incomes by only 0.021% relative to 15°C whereas the analogous loss of annual total income is 0.059%. Both of these features of the response, the smaller magnitude and the lower breakpoint temperature, are broadly consistent with the response of labor supply documented by Graff Zivin and Neidell (2014) (Figure 9D) and labor productivity responses from lab experiments (Seppanen, Fisk, and Lei (2006)). As with the crop yield response, the breakpoint documented by Graff Zivin and Neidell (2014) (∼25°C) is a higher temperature than what we observe in non-farm income (15°C). This difference is likely due in part to Graff Zivin and Neidell (2014) using daily maximum temperature rather than daily average temperature as we do—although the 10°C difference might be too large relative to normal diurnal temperature variations to be fully explained by this fact alone.[28] It is possible that changes in the quality of labor, i.e. the intensive margin, are responsible for this lower turning point: lab studies summarized in Seppanen, Fisk, and Lei (2006) indicate that productivity begins to decline at slightly lower temperatures (∼21-22°C). We observe that the point estimate for non-farm income increases in the hottest temperature bin. However, this point estimate is noisy and is neither statistically different from zero nor from the negative estimate at the adjacent temperature bin.

Quantitatively, 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. For a day with an average temperature of 25°C, annual non-farm income is estimated to fall by 0.000213 log points, which corresponds to a loss of 7.8% of an average 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

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

[28] The 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.

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.

Finally, we examine 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 share to our baseline specification in Equation 28 to identify a component of temperature-related income variation that projects systematically onto the spatial distribution of manufacturing. The results, shown in Figure 9F, indicate an inverted-U shaped relationship between manufacturing income and daily temperature that peaks around 9-12°C and declines at warmer temperatures, although this response is also somewhat noisy.

6.3 Evidence of specific substitutes for climate in production

Another strategy for understanding what mechanisms underly the marginal product of temperature is to examine how non-climatic allocation decisions embodied by $\mathbf{b}^*$ influence the observed marginal product. If populations adapt to certain temperature distributions by choosing one allocation and then adjust that allocation to compensate for a different temperature distribution while maximizing production, that indicates some factors in $\mathbf{b}$ are effective but costly substitutes[29] for certain dimensions of $\mathbf{C}$. Optimizing populations will essentially substitute human-made capital for "natural capital" (Solow, 1991). Here, natural capital can be conceptualized as the temperature distribution, with an implicit price equivalent to its marginal product. Next, we look for empirical evidence of such substitution by examining how the marginal product of temperature changes based on overall allocation of specific factors thought to be important for adaptation to temperature. In the graphical solution shown in Figure 2, this exercise can be thought of as examining the gradient of the value function when the sample of data is restricted to specific cross-sections $\mathbf{b} \in [b_1, b_2]$ in the $\mathcal{B}$-subspace.

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 in certain production processes is fairly obvious. 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 thought 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).

[29] Were these substitutions not costly we would likely observe similar allocations in all climates.

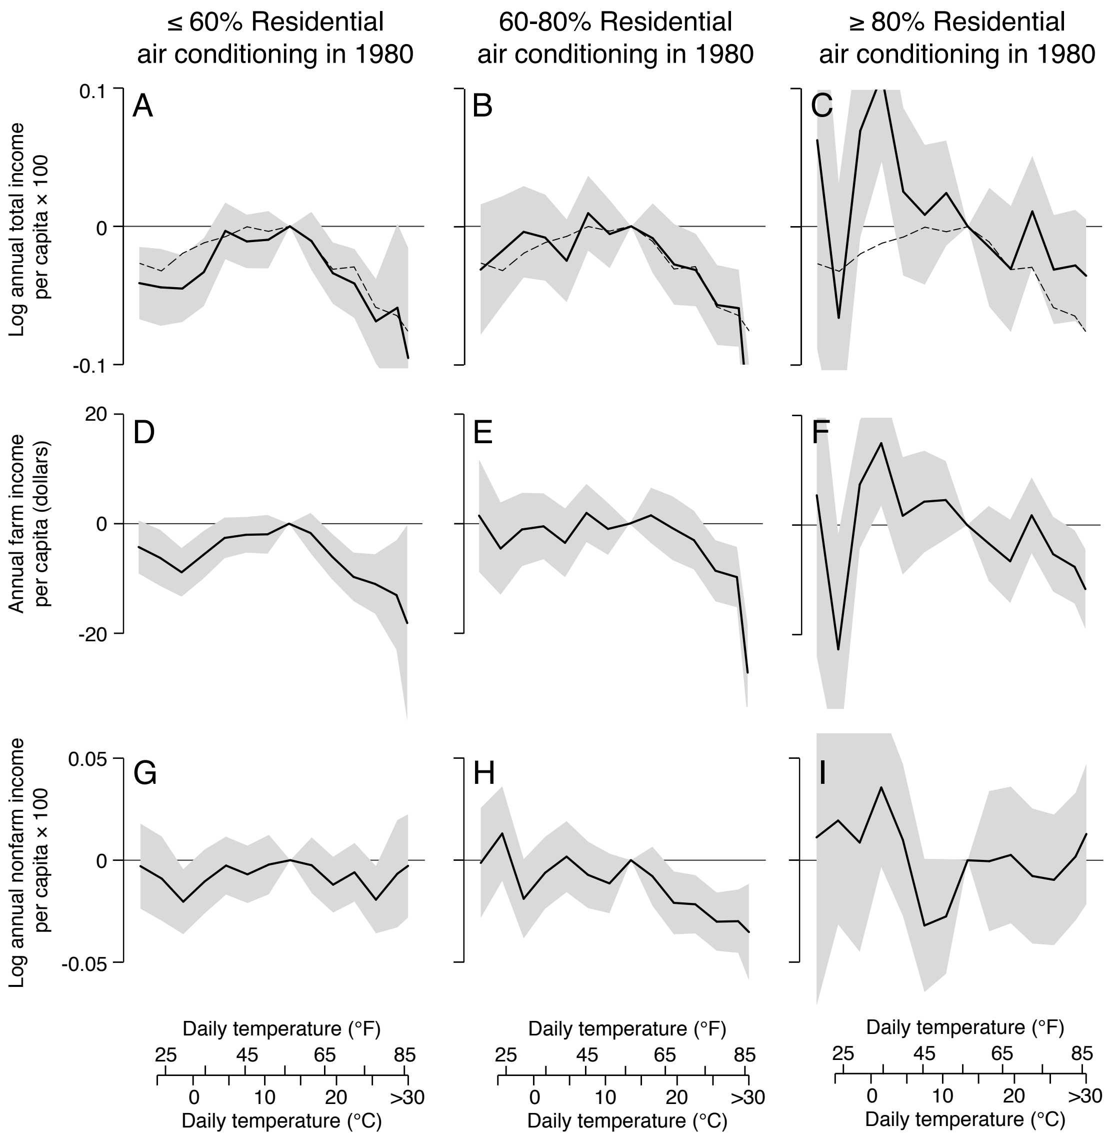

Figure 10: Estimated effect on income by 1980 residential air conditioning penetration, as reported by the 1980 Census. First row shows the marginal product of temperature for total income (Panels A-C, dashed lines correspond to the pooled estimate from Figure 4A); second row shows the effect of temperature on farm income (Panels D-F); third row shows effect of temperature on nonfarm income (Panels G-I). First column show results for counties that reported $\leq$ 60% air conditioning penetration; second column correspond to counties with penetration rates between 60 and 80%; third column corresponds to counties with penetration rates $\geq$ 80%. Results for all subsamples (within a single row) are jointly estimated. Shaded area is 95% confidence interval.

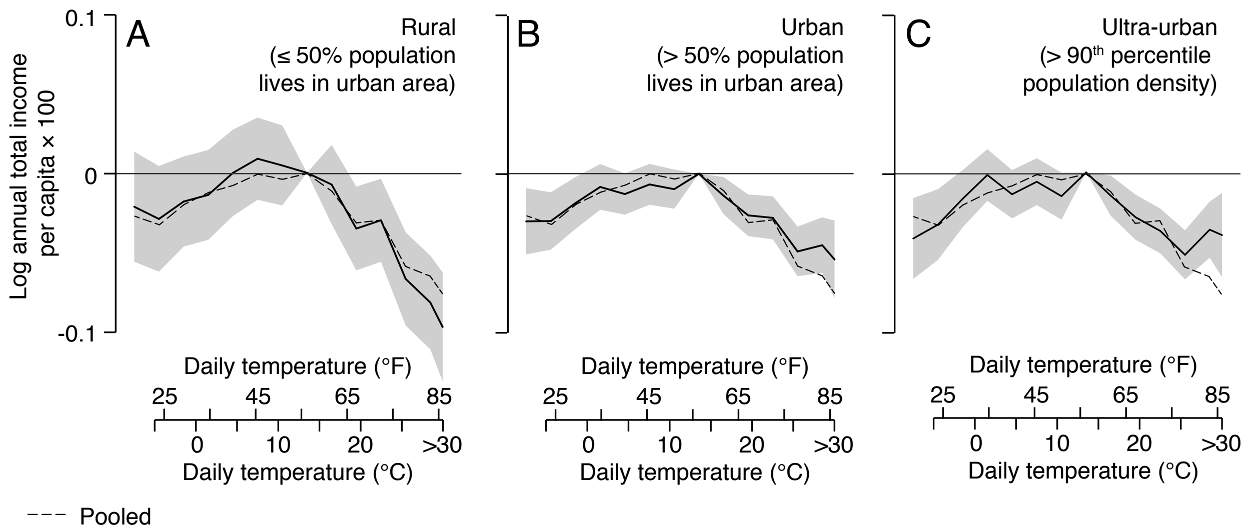

Figure 11: (A) Effect on income per capita for counties where ≤ 50% of residents lived in an urban area, as reported by the 2010 Census. (B) Effect for counties where > 50% of residents lived in an urban area. (C) Effect for counties in the 90th percentile of population density and above. Results for panels A and B are jointly estimated. Results for panel C are jointly estimated with counties below the 90th percentile of population density (not shown). Shaded area corresponds to the 95% confidence interval. Dashed lines correspond to the pooled estimate from Figure 4A.

Air conditioning We classify the counties in our sample into three groups based on the degree of residential air conditioning (AC) penetration reported in the 1980 Census: (1) counties with penetration rates less than or equal to 60%, (2) counties with penetration rates between 60% and 80%, and (3) counties with penetration rates of 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 (28) 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 whole sample (Panels A and B). Counties with the highest AC penetration rates, on the other hand, appear to be half as susceptible to such temperatures, and the point estimates for this group of counties are not statistically significant (Panel C).

The next six panels show separate results for farm and non-farm income. Non-farm income in counties with the highest rates of AC penetration is essentially immune to high temperatures (Panel I). By contrast, counties with AC penetration rates of 60%-80% experience increasing declines in non-farm income with higher temperatures (Panel H). Counties with the lowest rates of AC penetration fall somewhere in between, with point estimates that are negative but largely statistically insignificant (Panel G). Interestingly, farm income in counties with the highest rates of

AC penetration is also less susceptible to heat than counties with lower rates of AC penetration (Panel F vs. Panels D and E), suggesting that the former may also be more likely to adopt other adaptation strategies, such as a more heat-tolerant crop mix.

Urbanization Next, we classify the counties in our sample by whether a majority of their population lived in an urban area in 2010, as reported by the U.S. Census Bureau.[30] About 41% of U.S. counties are "urban" by this definition. Alternatively, we classify a county as urban if its population density averages at least in the 90th percentile during our sample period. By definition, only 10% of counties will be considered urban in this case. As with AC penetration, we interact the contemporaneous temperature and precipitation bins with urban/rural indicators to arrive at our estimates.

The results for total income per capita are shown in Figure 11. Rural counties are slightly more susceptible to the hottest temperatures than the entire sample, although the differences are not statistically significant (panel A). By contrast, urban areas exhibit a lower susceptibility to heat (panel B). Counties at the top of the population density distribution are even less affected by high temperatures, although they remain far from immune from them (panel C).

Overall, our results suggest that AC penetration is a much more important predictor of susceptibility of income to heat than is urbanization, although both clearly matter. Because allocations of resources towards AC installation and usage will likely be a direct response to warmer temperature distributions, the costs and benefits of AC allocations will already by implicitly captured by the "full adaptation model" that allows the value function to contain curvature (e.g. declining marginal effects of high temperatures seem likely to be a result of AC technology). It is less obvious that the influence of urbanization is explicitly captured in any dimension of our previously estimated value function. Rather, our main estimates simply reflect average treatment effects over the present joint-distribution of urban populations and temperature climates. Thus, in order to account for both the effects of AC (and other endogenous adaptations) as well as the influence of urbanization in our valuations of the current and future climate (below), we implement a "full adaptation" model of the value function that contains surfaces for both urban and non-urban counties. This is implemented by interacting an "urban" dummy variable with every polynomial for each temperature bin.

6.4 Other features of the economic response to temperatures

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 is unknown. To examine whether government transfers might be contributing to these results, 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 (see Appendix for details). We

[30]To our knowledge, this statistic is not available in earlier years.

find that daily temperatures have zero effect on county-level annual transfers from the government (excluding crop-related payments) or on county-level spending on public medical benefits (see Panels A and B of Appendix Figure A2). We find some evidence that *ad hoc* crop disaster payments increase as a results of really hot days (> 30°C), while crop insurance payouts increase steeply for days that exceed 27°C (80.6°F) (see Panels C and D of Appendix Figure A2). The latter estimates suggest that farm income losses would be roughly 25% higher if crop insurance were not available.

Spatial displacement Finally, we examine whether there is any evidence for spatial displacement of economic activity by estimating a spatial lag model with five 100-km distance bins. The results, shown in Appendix Figure A3, indicate that high temperatures continue to have a negative effect on own income even when accounting for neighbors' temperature. 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.

7 Valuing current and future climates

The robustness and stability in our estimates provide confidence that many county-level economies face a similar value function such that they lie on the same $V(\mathbf{C})$ surface up to a county-specific constant μ_i, which captures all other county-specific factors $\mathbf{z_{it}}$, and a mean zero idiosyncratic disturbance. Expected county income y when county i experiences its own current climate is then

$$E[y_i(\mathbf{C}_i)|\mathbf{z}_{it}] = V(\mathbf{C}_i) + \mu_i(\mathbf{z}_{it}). \tag{29}$$

By applying our estimates for the marginal product of climate $\widehat{\frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}}}$ at each point in the climate space $\mathcal{C}$ to Equation 26, we can trace out a change in income for i if the climate were displaced to a counterfactual climate $\mathbf{C}_{i2}$, net of all adaptive adjustments in $\mathbf{b}_i^*$, which we do not need to observe directly:

$$\begin{aligned} E[y_i(\mathbf{C}_{i2})|\mathbf{z}_{it}] &= V(\mathbf{C}_i) + (V(\mathbf{C}_{i2}) - V(\mathbf{C}_i)) + \mu_i(\mathbf{z}_{it}) && (30) \\ &= V(\mathbf{C}_i) + \underbrace{\int_{\mathbf{C}_i}^{\mathbf{C}_{i2}} \frac{\mathrm{d}\widehat{V(\mathbf{C})}}{\mathrm{d}\mathbf{C}}}_{\Delta y} + \mu_i(\mathbf{z}_{it}). && (31) \end{aligned}$$

We compute Δy for two counterfactual climates. First, we consider how much daily temperature distributions contribute to current production by taking a single county, displacing its climate to the climate of another county, and comparing how these counterfactual productivities differ. Second, we gradually distort the climate of each county along a "business as usual" climate change scenario and compute how its productivity changes through 2100. Both comparisons focus on the term Δy as they net out $V(\mathbf{C}_i) + \mu_i(\mathbf{z}_{it})$, in the first case by using a single benchmark county for all comparisons and in the second case by comparing a county to itself in later moments in time.

7.1 Contribution of temperature to current production

To understand how current temperature climates contribute to cross-sectional patterns of productivities, we integrate our fully non-linear estimate of $\widehat{\frac{\mathrm{d}V(\mathbf{C})}{\mathrm{d}\mathbf{C}}}$ from the benchmark historical climate (averaged over 1968-1990) of Lebanon, Kansas to the climate of each other county. We choose Lebanon simply because it is the geographic centroid of the country, aiding display of results, but comparisons of this Δy across locations are invariant to the constant of integrations, which in this case happens to be $V(\mathbf{C}_{Lebanon}) + \mu_{Lebanon}$.

The left panel in Figure 12 depicts the change in income that occurs as the climate of Lebanon (marked with a black circle) is smoothly transitioned to the climate of each county in our pooled "full adaptation" model. This panel depicts the value function evaluated at historical temperature climates across the country, net of all county-specific differences and all adaptive adjustments. The white band of counties, stretching through the corn belt, south of Appalachia, and up to the Mid-Atlantic states, indicates climates that are similar in 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, due to the high number of hot days (in the current year) in those climates. Note that these values fully account for the curvature in $V(\mathbf{C})$ shown in Figure 8, which captures 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 locations along the Pacific, Rockies, Great Lakes, and New England earning \$500-1000 or more per capita in annual income. Note that accounting for the dynamic effects of these earning losses would increase their magnitude by a factor of 5.01, as discussed above.

The right panel of of Figure 12 shows an analogous calculation, but, instead of using a fully pooled model, it conditions the value of the climate on whether a county is urban or not. Using this approach, regional patterns are largely unchanged, but urban locations exhibit a more muted version of regional patterns. This occurs because the marginal effects of climate appear smaller in overall magnitude in urban regions, as shown earlier, suggesting that the human-made capital installed in urban locations is partially a substitute for cool climatic conditions. Accounting for urban-rural heterogeneity in this way is akin to treating currently installed urban capital as part of each county's endowment that mediates the effect of the climate. The distribution of urban counties is not systematically correlated with climatic conditions, so this capital might be thought of as allocated according to an orthogonal optimization. The pooled model (left panel) can be considered the *a priori* valuation of each county's climate if the urban-rural status of each county were unknown in advance of the valuation calculations.

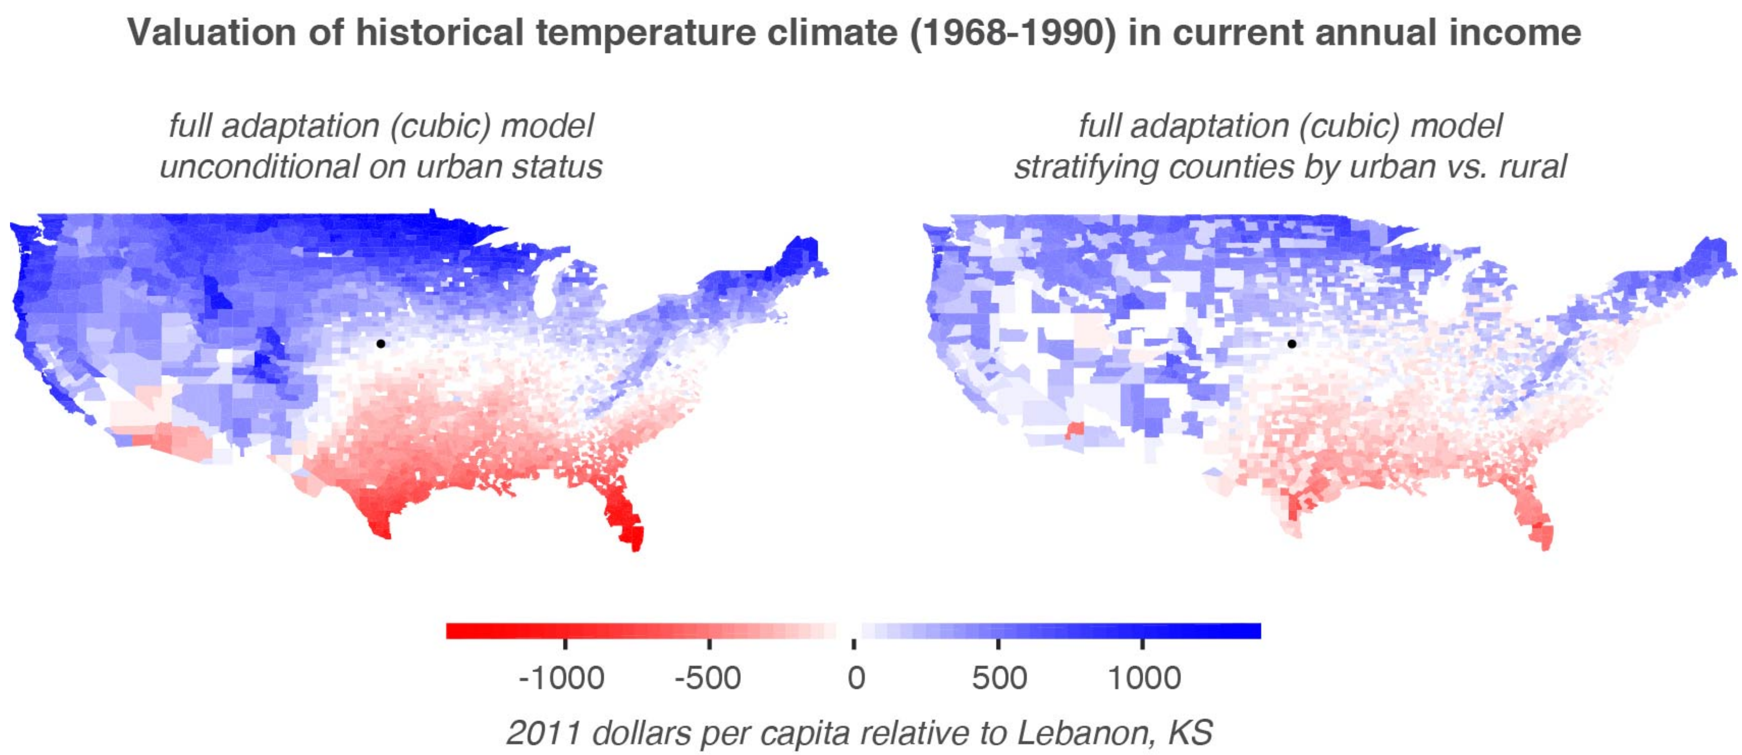

Figure 12: Surfaces describing the value of historical county climates in dollars of annual production relative to the climate of Lebanon, Kansas, the geographic centroid (black circle). Values are computed by integrating the marginal effect of distorting the climate of a county to that of a neighboring county, and so on, according to the empirically estimated marginal effects of these distortions. The left map uses the globally pooled model that accounts for adaptation by allowing the effect of temperatures in each of 16 bins to evolve as a unique third order polynomial in the count of days within each bin. The right map uses a model that is stratified by whether a county is urban or rural. Although these maps display cross-sectional differences, the empirical model used to construct these surfaces is estimated only using within-county variation in the distribution of daily temperatures.

7.2 Production distortions due to future warming

Next, we use 44 different climate change scenarios from Hsiang et al. (2017) to project how output will change due to future warming in RCP8.5 ("business as usual") relative to a counterfactual projection where temperatures remain at historical levels. Used as an ensemble, this set of projections is constructed to emulate the distribution of global climate sensitivities. For each U.S. 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. 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 or projected population, assuming that population growth in each county follows a linear trend.

We apply the warming projections to three sets of estimates: the affine model where all counties are pooled together and the full adaptation model (cubic in the number of days in each temperature bin) model, with and without distinguishing between rural and urban counties, as measured by whether the majority of a county's residents lived in an urban area in 2010. The affine model assumes that each additional day in a temperature bin has the same marginal effect on income, essentially ruling out any effects of adaptive adjustments that alter the marginal effect of additional warming. The cubic model, on the other hand, picks up non-linearities in a very flexible manner, allowing us to account for adaptation as a function of the number of days in each temperature bin (recall Figure 8). However, this model will not capture dimensions of adaptation that are not correlated with the temperature distribution, such as urbanization. For this reason, we also create projections using a full adaptation model that is fully interacted with an urban indicator—effectively estimating two curved value functions, one that applies to urban counties and one that applies to all other counties.

The spatial distribution of changes in the NPV of total income for the median climate trajectory (in terms of the total income loss) is shown in the first column of Figure 13. Dollar amounts correspond to billions of 2011 dollars. Here, we use a 3% discount rate for clarity, relegating discussion of other rates to Table 1. Without accounting for non-linearity or heterogeneity, the largest aggregate losses from climate change are 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, more counties, especially in Texas, are projected to experience income gains as a result of climate change; allowing for urban-rural heterogeneity produces more gains in urban counties in Gulf Coast states. Yet, as discussed above, 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

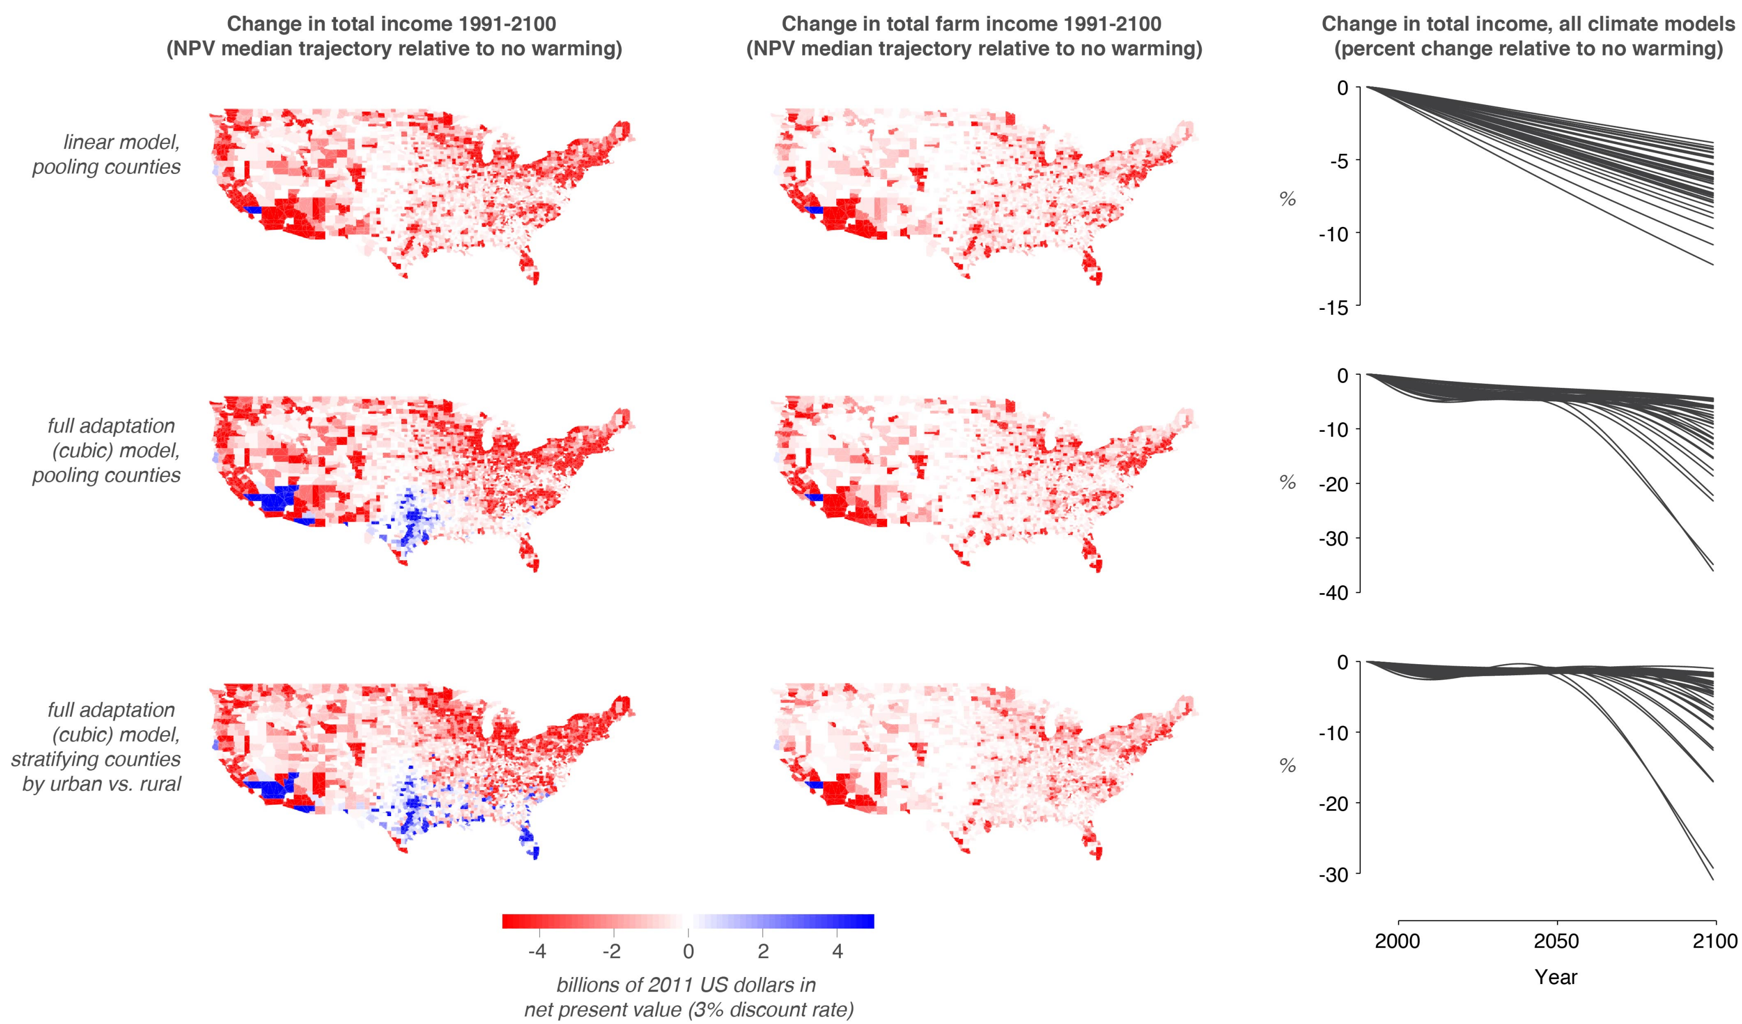

Figure 13: Outcome variable specified above each column. Projections in the rightmost column created by combining estimates from models specified by row headings with 44 model projections for 2080-2099 from Houser et al. (2015). Maps show the spatial distribution of the net present value of income losses for the median climate projection, as measured by the estimated impact in 2099. 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.

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: affine 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.

effects of Southern states are no longer pooled with these cooler locations (recall Figure 8). Because economic production is more heavily concentrated in the North, where allowing for curvature in the value function increases damages, than the South, where allowing for curvature decreases damages, allowing for adaptation by letting marginal damages vary by climate has the net effect of *increasing* total national income losses (shown in Table 1).

By contrast, projected changes in farm income (set of maps in the second column) are much more similar across the three sets of estimates. Very few counties are projected to gain agricultural income as a result of climate change. The largest losers are again concentrated in the Southwest, Northeast, and Florida. Comparing the two columns, we see that while farm income losses contribute to aggregate losses, they do not fully explain their magnitude, especially once heterogeneity and non-linearities are accounted for. Non-agricultural losses thus must play a major role in generating the bulk of projected reductions in economic output.

The third column of Figure 13 shows how aggregate income losses accumulate over time across the 44 climate projections. By definition, income losses are zero in 1990. In the linear model, aggregate losses grow linearly between 1991 and 2099, and range from 3.8% to 12.2% of the no-climate-change counterfactual annual income. The median scenario predicts income losses of 6.3% by 2099. Both projections based on cubic estimates (with and without allowing for urban-rural heterogeneity) display non-linear and non-monotonic temporal 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 urban-rural 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 (affine, homogenous cubic, and heterogeneous cubic) and for three different discount rates (1%, 3%, and 5%). Panel A shows the affine model estimates. With a 3% discount rate, the median NPV of income loss is $13.7 trillion in 2011 dollars. At the 10th percentile of the distribution, the losses are almost $19 trillion, and at the 90th percentile losses are estimated at $9.6 trillion. With respect to the discount rate, the NPV of losses is more than two times larger when we use a discount rate of 1% rather than 3%. Conversely, it is about one quarter to one third lower than the 3% estimates if we use an annual discount rate of 5%.

Allowing for a cubic relationship between income and the number of days in a temperature bin (Panel B) increases income loss estimates across the board. As discussed above, this occurs because the marginal effects of warming (for the high temperature bnis) are negatively correlated across space with the overall economic output of counties. However, adding urban-rural heterogeneity while still allowing for a cubic relationship (Panel C) reduces projected losses and yields the smallest income losses. This adjustment has substantial negative effects on 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 an economy influenced by its climate have struggled to simultaneously account for observable factors that differ across locations and the overall impact, net of costs, of constantly re-optimized adaptive adjustments. Here we developed a general reformulation of the problem that, when considering the role of the climate in economic production, delivers both. For a macro-economy in general equilibrium, market optimization leads to the maximization of certain aggregate quantities, such as total income. This maximization allows us to leverage the Envelope Theorem such that random perturbations in weather can be econometrically exploited to identify the local marginal product of climate. In cases where we have a large number of small macro-economies that are densely packed in climate-space (such as US counties), we may integrate many "nearby" estimates of these local marginal effects to recover the entire value function describing the total product of climate in these economies, which characterizes the economic impact of nonmarginal changes to the climate. Importantly, this value function captures the net effect–both costs and benefits–of all adaptive adjustments in the economy and, so long as the estimator used purges unobservable heterogeneity, it will not be biased by cross-sectional covariates unrelated to the climate.

Applying this approach to estimate the role of daily temperature distributions in modern US markets, we find that the marginal product of climate, which is the local gradient vector of the value function, is remarkably stationary over time and fairly stationary across space. We find that the gradient vector rotates as an economy's climate changes, the PPF adjusts, and the population reallocates resources in response, such that the marginal impact of very high temperatures is reduced. Some of this response appears due to compositional shifts in production away from agriculture and the usage of air conditioning. Urbanization, which does not appear directly related to climate, also plays a role in determining the local marginal product of climate. Accounting for curvature in the value function, which captures all adaptations, as well as urban-rural heterogeneity, we compute causal estimates for cross-sectional differences in income across US counties attributable to each county's climate by integrating each county's position on the value function relative to a benchmark county. To the best of our knowledge, this is the first such calculation of this type. Using a similar technique, we estimate distortions to national economic output induced by future anthropogenic warming in the RCP8.5 scenario across a range of climate models, fully accounting for both the costs and benefits of adaptation through factor reallocation. Importantly, counter to a widely invoked "folk theorem," accounting for adaptation by allowing for curvature in the value function increases forecast losses because marginal damages from warming are larger in Northern regions where economic output is concentrated, a fact that is obscured in affine models that assume the marginal product of climate is constant across space.

There are numerous important caveats for this analysis. First, these results depend on markets being essentially efficient in the long run. We assume agents have 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, it is more difficult to analyze the long-run general equilibrium structure of the economy and, critically, our Envelope Theorem result that depends on efficient market-clearing will no longer hold exactly. Future work may examine how market distortions, imperfect information, and the incomplete rationality of decision-makers may alter these findings.

Another key assumption in our analysis is that disturbances due to weather are "small" such that they do not move the economy "too far" from its equilibrium. This assumption is important for guaranteeing the identifiability of the marginal product of climate using weather variation, and its validity depends on the spatial scale of analysis, how weather and climate are defined, and how they relate to the economy. We chose to demonstrate the application of this result using annual distributions of daily temperature, described with temperature bins, in part because perturbing an annual temperature distribution by shifting one or two days from one bin to the next is plausibly a "small" perturbation in an otherwise large space of possible temperature distributions. This intuition appears to be confirmed by our results, since shifting individual days results in only fractional changes in annual percentages of income. Thus the "local" assumption we invoke seems valid in this context. Yet the validity of this assumption cannot be extended blindly to all measures of all dimensions of climate—for example, hurricanes and mega-droughts might not be sufficiently "small" economic perturbations for our core results to apply (Hornbeck, 2012; Deryugina, Kawano, and Levitt, forthcoming).

Perhaps most importantly, none of these results should be interpreted as welfare effects. The focus of this analysis has been to characterize the contribution of the climate to aggregate economic production that is captured by total market revenue. Many other analyses that consider the economic value of the climate in welfare terms also account for numerous non-market effects, such as the welfare loss associated with 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 climate but are not accounted for in the present analysis beyond any influence they have on the structure of the PPF or aggregate demand.

A core contribution of our analysis accounts for a large number of allocative adjustments made *within* a macroeconomy to cope with a change in the economy's climate. Our empirical implementation focuses on a large number of 'small' macroeconomies, US counties, within which a vast number of possible allocations for $\mathbf{b}$ exist. There is, however, interest in allocative adjustments *across larger spatial scales* that are responses to regional or global climatic changes (Desmet and Rossi-Hansberg, 2015; Costinot, Donaldson, and Smith, 2016; Dingel, Hsiang, and Meng, 2017; Desmet et al., 2017). Our theoretical analysis remains relevant to these larger scales, as the domain of the macroeconomy under consideration can simply be expanded to contain larger regions. However, our empirical analysis does not fully address many of the ideas explored in other studies.

Nonetheless, applying our empirical technique to larger units of analysis that span larger spatial scales is a valid approach to accounting for spatial reallocations across the larger regions of space contained within these units.[31]

When considering our projections of production distortion due to anthropogenic climate change, it is important to note that we only consider the impact of shifting distributions of daily temperature and we omit any influence of other climatic factors that may also change, such as hurricane frequencies or intensities, sea level rise, or drought. In principle, our approach can easily account for a large number of additional dimensions of the climate, something which is an important avenue for future work.

In addition, it is crucial to recognize that future, unknowable technological innovations that may eventually affect the marginal product of climate are unlikely to be well-captured in our empirical analysis. Theoretically, new technologies can been incorporated into our model by increasing the dimensionality of $\mathcal{B}$ to allow for allocations towards a new type of technology, which may have prohibitively high costs prior to its discovery. Because currently unknown technologies are by definition not available, it is not possible to empirically explore the structure of the value function in the subspace of $\mathcal{B}$ in which resources have been allocated towards a not-yet-existent technology. Nonetheless, future innovations may be captured by our analysis to the extent that they are represented in the present marketplace. For example, valuations of assets (e.g. stocks, bonds, farm land) that reflect market beliefs about the future trajectory of technology affect allocation decisions within the present market. Future 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 in future periods.

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: Journal of the Econometric Society* :265–290.

Auffhammer, Maximilian, Solomon Hsiang, Wolfram Schlenker, and Adam Sobel. 2013. "Using weather data and climate model output in economic analyses of climate change." *Review of Environmental Economics and Policy* 7 (2):181–198.

Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S Shapiro. 2013. "Adapting to climate change: the remarkable decline in the US temperature-mortality relationship over the 20th century." Tech. rep., National Bureau of Economic Research.

[31] In fact, the prior analysis by Burke, Hsiang, and Miguel (2015) essentially implemented just such an approach, albeit without articulating that it represented this exercise.

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 M 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).

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.

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.

Froom, Paul, Yeheskial Caine, Igal Shochat, and Joseph Ribak. 1993. "Heat Stress and Helicopter Pilot Errors." *Journal of Occupational and Environmental Medicine* 35:720–724.

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." *Journal of Environmental Economics 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 M. 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. 2016. "Climate econometrics." *Annual Review of Resource Economics* 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 M. 2010. "Temperatures and Cyclones strongly associated with economic production in the Caribbean and Central America." *Proceedings of the National Academy of Sciences* 107 (35):15367–15372.

Hsiang, Solomon M. and Daiju Narita. 2012. "Adaptation to Cyclone Risk: Evidence from the Global Cross-Section." *Climate Change Economics* 3 (2).

Jones, Benjamin and Benjamin Olken. 2010. "Climate shocks and exports." *American Economic Review: Papers and Proceedings* 100:454–459.

Kahn, Matthew E. 2013. *Climatopolis: how our cities will thrive in the hotter future*. Basic Books (AZ).

Kelly, David L, Charles D Kolstad, and Glenn T Mitchell. 2005. "Adjustment costs from environmental change." *Journal of Environmental Economics and Management* 50 (3):468–495.

Key, Nigel, Stacy Sneeringer, and David Marquardt. 2014. "Climate Change, Heat Stress, and US Dairy Production." USDA Economic Research Report Number 175.

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 Review* 84:753–771.

Nickell, Stephen. 1981. "Biases in dynamic models with fixed effects." *Econometrica* :1417–1426.

Nordhaus, William D and Zili Yang. 1996. "A regional dynamic general-equilibrium model of alternative climate-change strategies." *American Economic Review* 86 (4):741–765.

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. M. Hanemann, and A. C. Fisher. 2005. "Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach." *American Economic Review: Papers and Proceedings* 95 (1):395–406.

Schlenker, Wolfram, 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.

U.S. Census Bureau. 1969-2011. "County Business Patterns." Available from `https://www.census.gov/programs-surveys/cbp/data.html`.

Welch, Jarrod R., Jeffrey R. Vincent, Maximilian Auffhammer, Piedad F. Moya, Achim Dobermann, and David Dawe. 2010. "Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures." *Proceedings of the National Academy of Sciences* .

Appendix

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

Appendix Tables and Figures

Appendix Table A1: Summary statistics

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

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

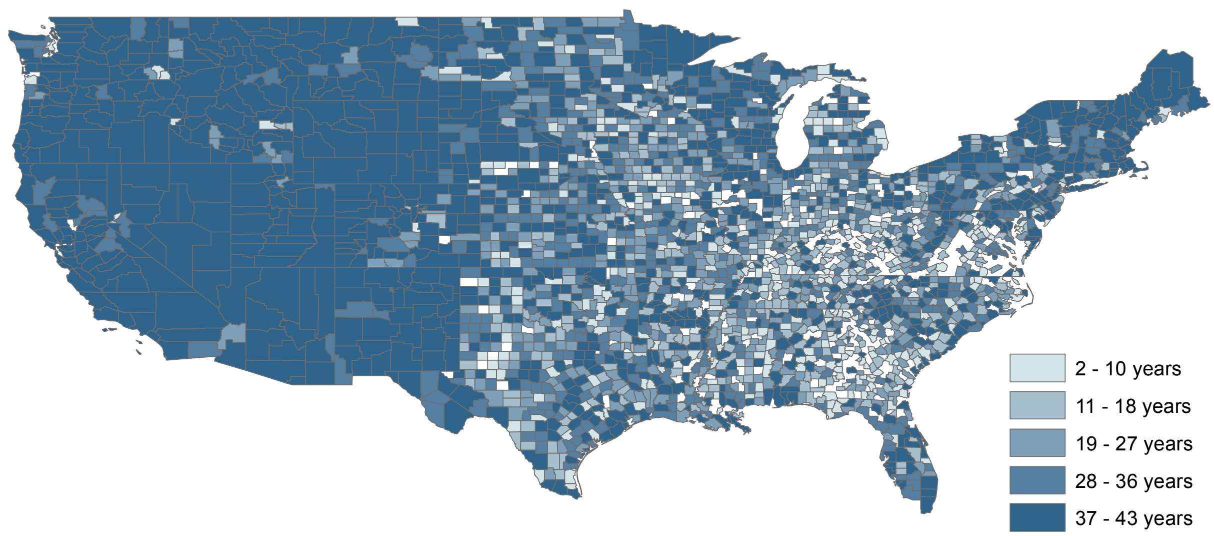

Appendix Figure A1: Number of years that each county has a complete record of daily average temperatures and daily rainfall. Years with incomplete records are dropped from the sample.

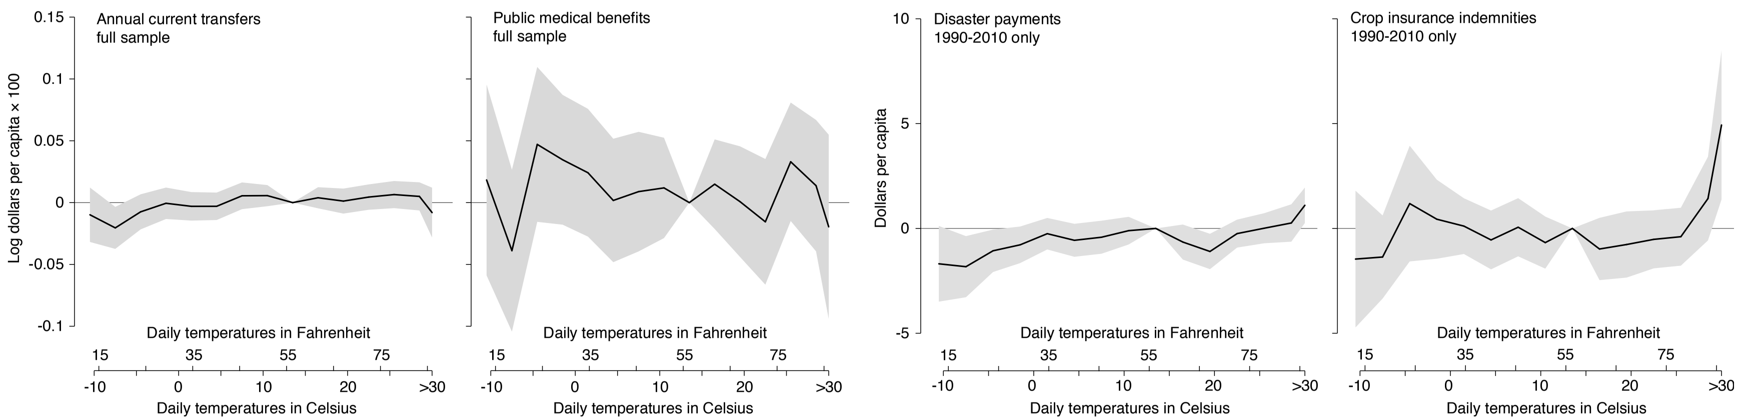

Appendix Figure A2: (A) The effect of daily average temperatures on log total government transfers to individuals per capita. (B) The same as A, but for the subset of transfers that are public medical benefits. (C) The effect of daily temperatures on *ad hoc* crop disaster payments per capita (in levels). (D) The effect of daily temperatures on crop insurance indemnities per capita.

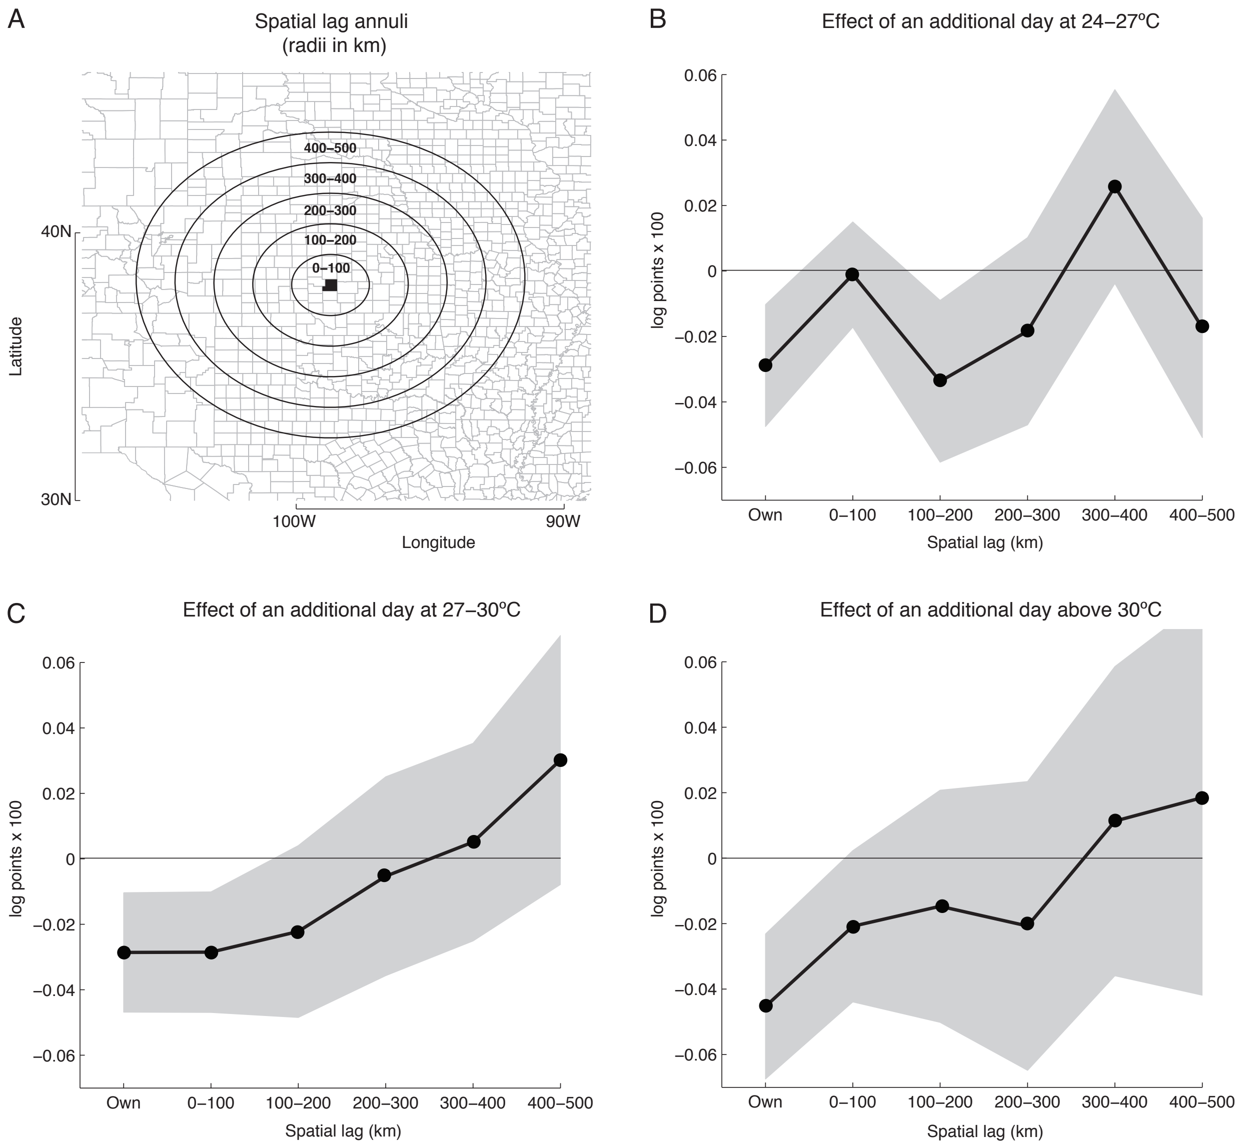

Appendix Figure A3: (A) Example of annuli used to construct spatial lags, relative to Stafford, Kansas (black). (B) Effect on i of each additional day at 24-27°C for i and 24-27°C days experienced by j's at various distances from i. (C) Same but for 27-30°C . (D) Same but for > 30°C . All effects in (B)-(D) are estimated simultaneously, along with own effects for lower temperatures and all controls.