

HEAT WAVES, CLIMATE CHANGE, AND ECONOMIC OUTPUT

Steve Miller

Environmental Studies Program,
University of Colorado Boulder

Jay Coggins

Department of Applied Economics,
University of Minnesota

Kenn Chua

Department of Applied Economics,
University of Minnesota

Hamid Mohtadi

Department of Economics, University of
Wisconsin-Milwaukee

Abstract

Climate change is likely to affect economies not only through warming, but also via an increase in prolonged extreme events like heat waves. However, the impacts of heat waves on economic output are not well captured by standard empirical approaches that ignore when hot days occur. Using a global dataset spanning 1979–2016, we show agricultural losses from past heat waves are up to an order of magnitude larger than suggested by standard approaches. Combining these estimates with a suite of climate models implies that by the end of the century, climate damages in agriculture may be 5–10 times larger than is predicted by a focus on mean temperature shifts alone. These findings have important implications for targeting and evaluating climate adaptation efforts. (JEL: Q54, Q51, O13)

1. Introduction

The costs and consequences of climate change are likely to extend well beyond the effects of gradual warming of average temperatures. In particular, tail events such as severe heat waves can prove fatal, reduce labor productivity and increase absenteeism, devastate crops, and strain power systems (Wahid et al. 2007; Kovats and Hajat 2008; Rocklöv and Forsberg 2008; Somanathan et al. 2015). The 2003 heat wave in France led to an estimated 11,000 deaths, the 2010 wave in Russia reduced grain yields by

The editor in charge of this paper was Bård Harstad.

Acknowledgments: We thank three anonymous reviewers, the editor, and seminar participants at the University of Minnesota, University of Wisconsin-Milwaukee, and AAEA and AERE conferences for helpful comments.

E-mail: steve.j.miller@colorado.edu (Miller); chuax025@umn.edu (Chua);
jcoggins@umn.edu (Coggins); mohtadi@uwm.edu (Mohtadi)

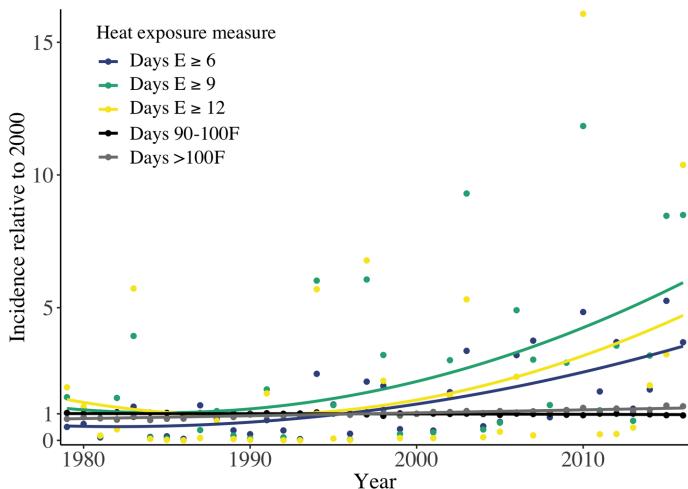


FIGURE 1. Trends in heat-wave incidence. Global average measures of exposure to heat relative to 2000 levels for example heat-wave metrics (in color) and counts of days in hot temperature bins (grayscale). Different colors reflect counts of days at or above different minimum exposure levels for the binary heat accumulation model (see Section 2.3 for details). Heat-wave measures are based on consecutive days of more than 1.5 standard deviations above average for a grid cell and time of year. Measures are aggregated to the country level using cropland and pasture weighting, then averaged across countries. Points represent actual global averages, while lines are quadratic trends fit to each series.

25% (\$15 billion), and widespread power outages and hospitalizations followed 2018 events in California, Japan, and Pakistan. Concerningly, heat waves are expected to become more common in the future (Meehl and Tebaldi 2004; Cowan et al. 2014; Russo et al. 2014; Mann et al. 2018; Coumou et al. 2018). Those projections match recent trends in which the incidence of prolonged exposure to heat is rising much faster than gradual shifts in the temperature distribution would suggest (Figure 1). Together, large damages from and rising frequency of heat waves suggest a full accounting of climate impacts requires more careful attention to prolonged exposure to heat.¹

Despite their apparently large impacts, heat waves have received insufficient attention from economists. A large and growing related literature examines the impacts of temperature on a range of outcomes, often linking the results to implications of climate change (see Auffhammer and Mansur 2014; Dell, Jones, and Olken 2014; Carleton and Hsiang 2016; Hsiang 2016, for summaries). Many of these studies examine the marginal effects of exposure to a single day of high temperatures, either on contemporaneous outcomes (e.g. Graff Zivin and Neidell 2014) or on outcomes for a month or year in which the hot day occurred (e.g. Schlenker, Hanemann, and Fisher 2006; Schlenker

1. In crops as in humans, the impact of heat depends on both temperature and duration of exposure (Wahid et al. 2007). For instance, heat shock proteins produced by plants in response to high temperatures can confer tolerance (Bita and Gerats 2013), but that tolerance decays with time (Charng et al. 2007); the intensity and timing of heat exposure matters.

and Roberts 2009; Barreca et al. 2016). Other researchers instead examine the impacts of shifts in monthly or annual mean temperatures (e.g. Burke, Hsiang, and Miguel 2015). Most integrated assessment models used to generate social cost of carbon (SCC) estimates share this focus on averages, considering only damages that arise from shift in mean temperatures (National Academies of Sciences and Medicine 2017). While valuable, neither approach isolates the multi-day exposure to abnormal heat which characterizes a wave. Instead, both approaches collapse a sequence of temperatures into one or a handful of summary statistics, discarding the time structure of heat exposure.

In response, we propose new measures of prolonged exposure to heat, estimate the impacts of past heat waves, and use those estimates to project future damages under climate change. Our metrics are straightforward, tracking prolonged exposure to temperatures that are abnormally high for a location and time of year and summarizing the number of days in which resulting heat stress falls in particular ranges (Section 2). Our core estimates, which come from a country-by-year panel spanning 1979–2016, reveal strong negative effects of heat waves on agriculture. For example, an abnormally hot day preceded by at least two others reduces agricultural output by approximately 0.3%. Higher levels of prolonged heat exposure are more damaging: after 8 days of abnormally high temperatures, each additional hot day reduces agricultural output by 1.7%.² Effects on non-agricultural and overall output are an order of magnitude smaller in percentage point terms, but because average output outside of agriculture is substantially larger (by a factor of 44 in our sample), those effects may still be economically significant. Taken together, our estimates imply that three high-profile heat waves (France 2003, Russia 2010, United States 2012) led to total losses of \$52 billion.

Looking ahead, our estimates raise concerns about potential impacts of climate change, particularly for agriculture. Applied to a suite of climate models for Representative Concentration Pathway (RCP) 4.5, our estimates suggest that without further adaptation, potential additional agricultural output losses due to heat waves could surpass 10% per year by the end of the century (Section 6.1.2). As a result, SCC estimates, which typically ignore damages from extremes, are likely to be too low. Adaptation could, of course, dampen impacts from heat waves and mitigate bias in SCC estimates. If the temperature thresholds underlying our heat wave measure change at historical rates, reflecting gradual adaptation to new temperature regimes, projected agricultural damages from heat waves fall to a country-level average of 4.5% per year. While smaller, those losses remain substantially larger than the projected 0.9% losses arising from mean temperature shifts. Moreover, a retrospective assessment of adaptation paints a less encouraging picture: we find no evidence that impacts of heat waves have declined through time in agriculture (Section 6.3.2 and Online Appendix Table A.8). The rarity of major heat waves poses challenges for

2. Estimates reflect heat waves of consecutive days with temperatures at least 1.5 standard deviations above the local average during the historically hottest 90 days of the year.

selecting optimal levels of adaptation (see Section 7), even if adaptation to heat waves and gradual warming entail similar types of protective investment.

The importance of explicitly considering prolonged exposure to heat is made clear by comparing our metrics and associated estimates to existing approaches. Comparisons of our approach with standard measures of temperature exposure (e.g. annual mean and counts of days in fixed temperature bins) reveal low to modest correlations (Online Appendix Table A.1), suggesting our heat-wave measures provide new information not present in standard measures of heat exposure. The new measures also successfully pinpoint major heat-wave events, such as those in France in 2003 and Russia in 2010 (Figure 4), while conventional measures do not (Online Appendix Figure A.6). More importantly, we find the effects of heat waves to be substantially larger than what is implied by binned temperature measures. For example, our point estimates suggest the 2003 heat wave in France led to \$3.1 billion in lost agricultural output—10 times larger than estimates using a standard approach counting days in ten-degree temperature bins (Table 5). A hybrid approach using our location-specific threshold for abnormal heat but ignoring whether or not hot days occur in a row still yields estimates less than one third the size of our approach (Table 6). More generally, total estimated damages from the three high-profile heat waves we consider are substantially smaller when using a standard temperature-bin approach (\$13 billion vs. \$52 billion). In short, accounting for prolonged exposure to high temperatures during heat waves yields meaningful quantitative differences in estimated damages. The fact that our estimates are meaningfully larger than those suggested by existing approaches may be attributable to the non-additive effects of prolonged exposure to heat: we document statistically different effects of a hot day at the start of versus later in a heat wave (Table 7). Moreover, we find the effect of heat waves is distinct from that of average temperature; both general exposure to and persistence of heat matter.

To be clear, some existing economic studies do allow for interactive effects of temperature across days; however, most focus on much longer time scales than those relevant for the heat waves with which we are concerned. Models with quadratic effects of annual mean temperature (e.g. Burke, Hsiang, and Miguel 2015) allow the effects of a hot day to depend on prior exposure, but impose equal dependency regardless of when the earlier high temperatures occurred. Others implicitly or explicitly allow the effects of heat to depend on prior temperatures, but only in a very general sense. For example, some researchers allow for heterogeneity based on whether heat occurs early or late in a typically hot season (Graff Zivin and Neidell 2014; Schlenker and Roberts 2009) or in a historically hot location (Carleton 2017). We complement and build on these efforts in two ways. First, we consider interactions at much shorter time scales, asking whether the effects of extreme heat depend on temperatures in the preceding several days, rather than earlier in the year or in prior years. Second, we account for the interactions investigated in those prior studies. Our thresholds for what constitutes abnormal heat (Section 2.3) allow for both within-season acclimatization and long-run adaptation via day and location specificity, respectively. In addition, we test for evidence of both long-run adaptation and medium-term interaction between heat waves and overall temperatures within the same year (Section 6.3.2).

We also formalize, extend, and complement prior studies that do focus explicitly on heat waves, most of which fall outside of economics. Some indirectly capture prolonged exposure through multi-day temperature averages (e.g. Nairn and Fawcett 2014; Somanathan et al. 2015). However, many temperature sequences could produce the same average but have different implications. To account for sequencing, other researchers count uninterrupted stretches of days that lie above a heat threshold (e.g. Meehl and Tebaldi 2004; Gasparini and Armstrong 2011; Massetti and Mendelsohn 2015). We define a family of heat-wave metrics that encompasses and generalizes that approach in several ways. First, we motivate our metrics using a simple theoretical framework linking production and the accumulation of heat exposure. Second, our empirical results emphasize two specific accumulation models: one in which heat accumulation depends only on temperature surpassing a threshold (“binary” accumulation), and one in which the intensity of the heat above that threshold also plays a role (“intensity-based” accumulation). The novel “intensity-based” accumulation model allows for both duration and intensity to contribute to heat stress.³ Finally, the accumulation models underlying our metrics reveal opportunities for formal comparison among measures of temperature exposure (Section 6.2).

While we focus on heat waves, our general approach may prove useful for examining prolonged exposure to other environmental stressors. Studies of economic impacts of pollutants (Graff Zivin and Neidell 2012; Heyes, Neidell, and Saberian 2016) often allow for lagged effects but rarely consider potentially non-additive effects of prolonged exposure. Accounting for prolonged exposure may matter less for some pollutants (e.g. ozone), but merits investigation when physical recovery from exposure can span multiple days (e.g. PM_{2.5}).

The remainder of this paper is organized as follows. We first motivate the potential importance of heat waves in a simple analytical framework. Next, we describe the weather and economic data we use. We then detail our proposal for the measurement of heat waves and compare our proposed family of metrics to common measures of heat exposure used in the economics literature. After establishing differences in the information captured by these metrics, we examine the effects of heat waves on overall, non-agricultural, and agricultural economic output. The final section provides a discussion and concluding remarks.

2. Quantifying Exposure to Heat Waves

2.1. Background

Heat waves are often characterized by two defining features: abnormally warm conditions and prolonged exposure. The first feature of heat waves, abnormally high temperatures, hints that identical absolute temperatures may have different effects if they occur

3. In Section 6.3.4, we also let cooler days provide incomplete relief from heat stress.

in different locations or at different times of year. A week of 90-degree weather early in a growing season may have a different effect than a week of the same temperature when crops are mature. Similarly, a 90-degree week in Dubai is likely to have a different effect than an identical week in Vancouver. While many economic studies use absolute temperature measures that do not directly reflect the notion of abnormality, standard modeling techniques can account for heterogeneity in both levels and effects of temperature. Location dummies can account for static differences in local climate, while interaction terms between short-term temperature exposure with local climate variables can get at differential effects of exposure to the same absolute temperature. While most studies allow for static differences across locations, many do not allow for the effects of heat to vary, which is defensible if the study is of limited geographic scope. For global analyses like ours, however, differences in local climate must be taken into account.

The second feature of heat waves, prolonged exposure, distinguishes a heat wave from hot days that occur apart. Cumulative exposure of people, crops, or some engineered systems to heat over several consecutive days may have detrimental effects on health and labor, survival and growth, or performance of infrastructure. Put differently, seven hot days spread throughout a summer may have very different effects than a week of uninterrupted heat.

Most temperature-exposure metrics used in economic studies do not directly measure prolonged exposure. Figure 2 illustrates this issue, displaying daily maximum temperatures in the grid cell containing Paris in 2003 (top panel) and 2015 (middle panel), as well as Tehran in 2003 (bottom panel). For context, Paris and much of Western Europe experienced a major heat wave during the summer of 2003, but not 2015. The left plots show the raw temperature series in each year, together with a dashed line indicating the annual mean. The right plots illustrate the typical binning approach described above, counting days that fall in various ten-degree bins. Using mean temperature discards both the within-year distribution of temperatures and the temporal structure of heat exposure. The binning approach retains variation by constructing an approximation to the within-year distribution of temperature, but still discards the temporal structure.⁴ The sustained heat during the major European heat wave in the late summer of 2003 stands out in the time series, but is obscured when examining distributions and is entirely absent when examining annual average temperatures (note the mean for 2015 is higher than that for 2003). In addition, comparing temperatures in Paris and Tehran in 2003 highlights the importance of locally appropriate definitions of abnormal heat. Temperatures in 2003 were regularly hotter in Tehran than in Paris, but due in part to adaptation, that year made no headlines for widespread health or other damages in Tehran.

Outside of economics, a number of heat-wave metrics have been introduced to directly incorporate the notion of prolonged exposure. Some simply average temperatures over shorter time scales, e.g. a 3-day period (Meehl and Tebaldi 2004).

4. The loss of temporal structure via averaging or binning likely matters less if considering a shorter period, for example, a week rather than a year.

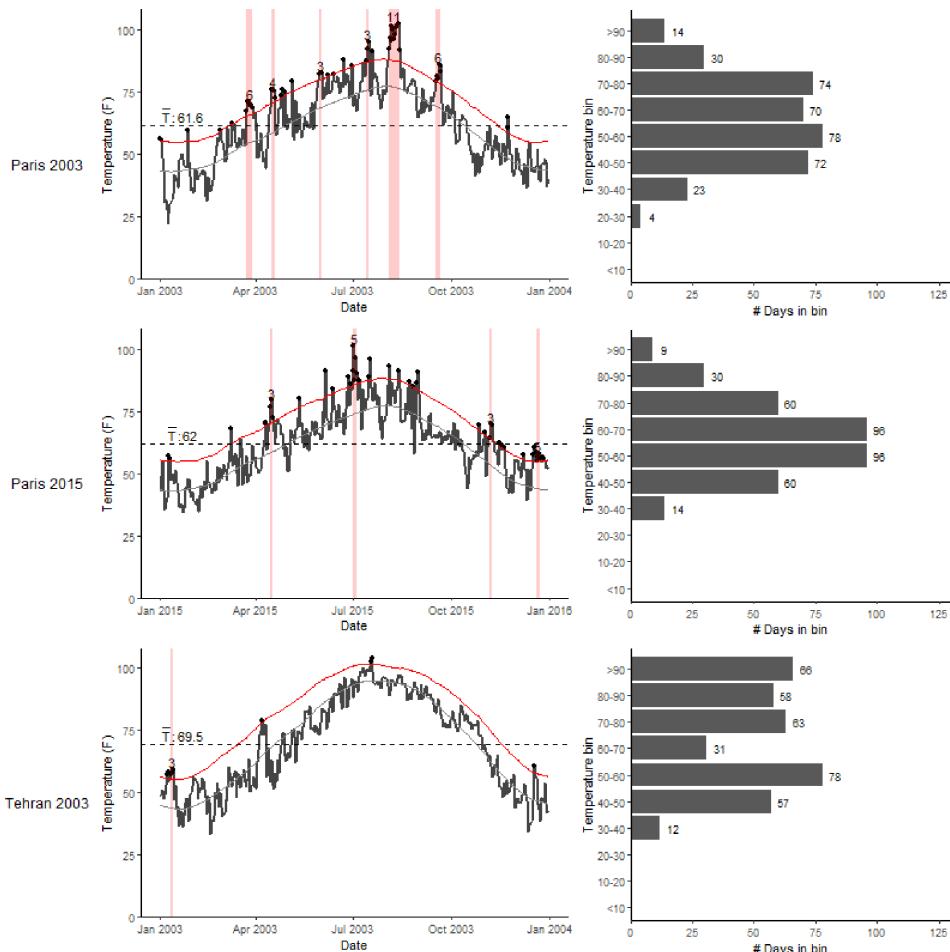


FIGURE 2. Existing and new temperature metrics. Temperature series (left), and bin-based approximation to the temperature distribution (right) for the grid cell containing Paris in 2003 (top), Paris in 2015 (middle), and Tehran in 2003 (bottom). Temperature series plots include daily cell temperature (black), long-run average per day (gray), and 1.5 standard deviations above average per day (red). Light red bands indicate heat waves using the simplest measure falling in our family of metrics: three or more consecutive days with temperature at least 1.5 standard deviations above average. Numeric annotations mark the duration of each heat wave in days.

Those multi-day averages can also be compared to longer-term averages as a way of incorporating the notion of abnormality (Nairn and Fawcett 2014) directly in a heat-wave index. More complex metrics combine multiple, quantile-based thresholds with minimum numbers of days that minimum temperature, maximum temperature, or both must exceed those thresholds (Meehl and Tebaldi 2004; Das and Smith 2012).

Such alternative measures more closely match the popular understanding of heat waves as prolonged exposure to abnormally hot temperatures, but are not without their own limitations. Potential concerns with average-based metrics, even over short

periods, were described above. The threshold-based metrics in Meehl and Tebaldi (2004) avoid those concerns; however, simply counting days above thresholds ignores potentially valuable information about *how far* above the threshold the temperature rises. In what follows, we propose a flexible family of heat-wave metrics to address these issues, incorporating both intensity and prolonged exposure.

2.2. Analytical Framework

To motivate our proposed measurement of heat waves and their effect on economic growth, we use a simple theoretical framework that builds upon Deryugina and Hsiang (2014) and Burke, Hsiang, and Miguel (2015). Our stylized model embeds two key features: both current and past temperatures can affect economic output on a given day, and those effects may interact. As a result, a serially correlated sequence of hot days may have a very different effect than the same number of hot days spread throughout a season or year. Further, the effects of a recent hot day on current output may depend on whether the intervening days were also hot.

We begin with a temperature-dependent Cobb–Douglas model of output Y_{yd} in an industry and location on day d of year y . For simplicity, we assume as in Burke, Hsiang, and Miguel (2015) that temperature only affects total factor productivity⁵ (TFP); capital K_{yd} , labor L_{yd} , prices p , and output elasticity of capital α do not respond. Letting the temperature on day d be T_{yd} and the sequence of temperatures up to and including day d be $T_{yd} \equiv [T_{y0}, \dots, T_{yd}]$, we write TFP as $f(T_{yd}, d)$ and output as $Y_{yd} = pf(T_{yd}, d)K_{yd}^\alpha L_{yd}^{1-\alpha}$.⁶ While the vector representation for T_{yd} does not directly measure heat waves, it does pave the way for introducing heat waves below by capturing serial correlation among the elements of T_{yd} . Further, letting TFP depend on the day d permits heat waves to have different effects at different times of year. For example, a heat wave is likely to matter much more for agriculture if it occurs during a growing season.

For context, we briefly note what assumptions common empirical analyses make about $f(T_{yd}, d)$. Many standard models assume that past temperatures do not matter for current production ($f(T_{yd}, d) = \tilde{f}(T_{yd}, d)$), and hence serial correlation in exposure to heat is irrelevant. This is true of models in which daily or longer-term average temperature enters linearly, as well as models using counts of days in various temperature bins to approximate a temperature distribution. Distributed lags of temperatures permit a role for serial correlation, but imply the effects of heat are additively separable across days, that is, $f(T_{yd}, d) = \sum_{d'=0}^d \tilde{f}_{d-d'}(T_{yd'}, d')$. Conversely, other specifications allow for non-separability, but again preclude effects of serial correlation. Adding quadratic or other nonlinear functions of mean temperature

5. In the Online Appendix (Section A.6), we relax this assumption by building upon Dell, Jones, and Olken (2012).

6. Implicit in this framework is the assumption that productivity depends only on temperatures in that location and not temperatures elsewhere.

requires any interdependency in temperature impacts to be the same whether two hot days occur back to back or several months apart. We find these restrictions to conflict with intuition about how prolonged exposure to heat is likely to operate.

To address these limitations, we assume TFP depends upon temperature as follows:

$$\begin{aligned} f(T_{yd}, d) &\equiv f^h(E_{yd}, d) + f^c(T_{yd}, d) \\ E_{yd} &\equiv E(T_{yd}). \end{aligned} \quad (1)$$

This formulation embeds two key assumptions. The first is that TFP can be decomposed into two terms: one (f^h) that depends upon the history of temperatures through day d , and a second (f^c) that depends only on the current temperature. The first term is the focus of our analysis, as it captures potential effects of prolonged exposure to heat. Our second key assumption in equation (1) is that effects of prolonged heat exposure on production operate through a scalar-valued summary function E . The resulting scalar is denoted by E_{yd} , and we refer to it as our prolonged heat-exposure index. In the next section, we discuss the particular functional form of E we adopt for our empirical application.

Our summary of prolonged heat-exposure using a model-based scalar is of course not the only option. A natural extension of the standard temperature-bin approach would slot the temperatures on the most recent D days into one of B bins, and define E_{yd} to be the length- D sequence of those bin identifiers. For example, if $B = 2$ and $D = 6$, a sequence 000111 might indicate that the most recent three days were abnormally hot (1), while the three before that were not (0). Because there are B^D such sequences, that approach grows empirically unwieldy beyond modest values of B and D . Still, in Section 6.3.4, we use this approach with $B = 2$ and $D = 6$ to provide evidence that temperature sequencing matters for output. We also show how our scalar model for E_{yd} is a special case of this approach, with the advantages of greatly reduced collinearity and improved interpretability.

To arrive at total output in a country and year we must aggregate production across days, locations, and industries. If TFP takes the form in equation (1) and production is exchangeable across time and space as in Burke, Hsiang, and Miguel (2015), we can write total production in a year as follows:

$$Y_y = \int_{-\infty}^{\infty} g^E(\tilde{E}) f^h(\tilde{E}) d\tilde{E} + \int_{-\infty}^{\infty} g^T(T - \bar{T}_y) f^c(T) dT. \quad (2)$$

Here g^E and g^T are densities describing exposure to specified levels of prolonged heat exposure and temperature, respectively. The second term in equation (2) is the subject of Burke, Hsiang, and Miguel (2015), which studies the nonlinear effects of average temperature on production.

The first term in equation (2) is the focus of our paper. It says that annual output is impacted by temperatures in a way that is not fully captured by the mean. Specifically, it allows for prolonged exposure to heat to influence output. The effect is simply a weighted sum of the components of TFP pertaining to prolonged heat exposure, with weights determined by the incidence of different levels of prolonged heat exposure.

Importantly, we allow f^h to be negative. Many high-profile damages from heat waves result from the destruction of accumulated output (e.g. crops). In our model, this is captured by $f^h(E_{yd})$ turning negative for sufficiently large E_{yd} . During the brief periods when E_{yd} does grow that large, prolonged exposure to heat may not only temporarily prevent positive contributions to annual output, but may also more than offset productivity on other, cooler days.

If the incidence of prolonged exposure to heat (g^E) could be measured, equation (2) could be estimated to learn about the component f^h of TFP. Operationalizing these ideas requires defining cumulative exposure E_{yd} , which we turn to next.

2.3. Measuring Prolonged Exposure to Heat

If prolonged exposure to heat does impact productivity, how should that exposure be measured? A metric could be constructed using knowledge from other fields about the way in which biological or engineered systems experience prolonged exposure to heat, potentially embedding many nonlinearities. However, since people, crops, and power systems are likely to respond to and accumulate heat in substantively different ways, designing a metric to reflect any one system may obscure relationships of interest for others. We thus seek a system-agnostic definition of prolonged exposure for the study of aggregate outcomes and for comparison of output across different economic sectors. We derive our measure based on a simple model of heat accumulation.

We define prolonged heat exposure E_{lyd} for a given day d in year y at location l as an evolving index with intentionally simple dynamics: it goes up on days hotter than a location- and day-specific threshold chosen to reflect “abnormal” heat, and it declines on other days, subject to censorship at zero. Events that most people would recognize as a heat wave, with several abnormally hot days in a row, will thus drive the index above zero for several days in a row. The index should reach its highest levels during those waves. In contrast, during periods of normal or abnormally cool weather the index will remain at zero, while sporadic hot days will push the index only briefly and moderately above zero.

Specifically, we define a family of indices E_{lyd} as follows:

$$\begin{aligned} E_{ly0} &= 0 \\ E_{lyd} &= \max \left\{ 0, E_{lyd-1} + 1(T_{lyd} \geq \underline{T}_{ld}) h_+(T_{lyd}, \underline{T}_{ld}) \right. \\ &\quad \left. - 1(T_{lyd} < \underline{T}_{ld}) h_-(T_{lyd}, \underline{T}_{ld}) \right\}. \end{aligned} \tag{3}$$

Heat exposure accumulates via a function h_+ when temperatures rise above the “abnormally warm” threshold \underline{T}_{ld} , and relief is provided via a function h_- when temperatures fall below \underline{T}_{ld} . We return to the choice of \underline{T}_{ld} later.

Here we focus on two specific choices for the dynamics in equation (3). To facilitate both exposition and comparison with prior work (Section 6.2), our main results use a “binary” accumulation model. In that model, all that matters is whether or not the

temperature is above a critical threshold, and a single day below provides complete relief from any heat experienced recently ($h_+(T_{lyd}, \underline{T}_{ld}) = 1$ and $h_-(T_{lyd}, \underline{T}_{ld}) = \infty$). In this case, the exposure index simply tracks the number of consecutive days the temperature has been above the threshold as of day d , which is either zero or the duration of an ongoing heat wave. The binary model thus underlies the common understanding of a heat wave as a stretch of uninterrupted hot days measured by its duration.

We also present complementary results (Section 6.3.1) using a novel “intensity-based” model, in which the accumulation of heat stress also depends on *how far above* the “abnormally hot” threshold a day’s temperature is. As a result, a heat wave with days 30F above normal will yield higher exposure indices E_{lyd}] than an equally long wave with days 20F above normal. In particular, we allow the accumulation of heat to follow a linear process: $h_+(T_{lyd}, \underline{T}_{ld}) = T_{lyd} - \underline{T}_{ld}$.⁷ Conversely, a single day below the threshold provides complete relief from heat stress ($h_-(T_{lyd}, \underline{T}_{ld}) = \infty$), though we relax that assumption in Section 6.3.4. In sum, the intensity-based index E_{lyd}] allows both prolonged exposure and severity of temperatures to substitute in the production of heat stress.

A key innovation we offer is in how we summarize a chosen exposure index for a period (e.g. year) and region (e.g. country) of interest. Based on equation (2), and by direct analogy to the common practice of counting days for which temperature falls in particular bins, we propose to count days during which E_{lyd}] falls in distinct bins. Spatially weighting the resulting histograms based on the density of relevant economic activity (as measured by, e.g. population density or the fraction of a location devoted to agriculture) yields a step function approximation to the sum of g_i^E across industries. Formally, define the lower and upper thresholds for prolonged heat exposure bin b as $E_b^>$ and $E_b^≤$, respectively. Then our measure of interest at the country and year level is a vector W_{cy}] defined as follows:

$$W_{cy} \equiv \sum_{l \in c} \omega_{lc} W_{ly} \quad (4)$$

where ω_{lc} are spatial weights with $\sum_{l \in c} \omega_{lc} = 1$ and $W_{ly} \equiv [W_{ly1}, \dots, W_{lyb}, \dots, W_{lyB}]$ is a vector with b th element

$$W_{lyb} = \sum_{d \in D} 1(E_b^> < E_{lyd} \leq E_b^≤). \quad (5)$$

In a regression of output on W_{cy}], the coefficient associated with the b th element of W_{cy}] provides an estimate of the differential effect on output of an additional day falling in prolonged heat-exposure bin b instead of a reference bin. Defining the reference bin to contain only $E = 0$ provides a clean comparison to a day with no abnormal heat exposure. These coefficients can still be used to quantify impacts of particular heat-wave events. A counterfactual in which a wave did not occur implies

7. As we detail later, while the dynamics of the index are linear, its effects are allowed to be nonlinear (equation 5).

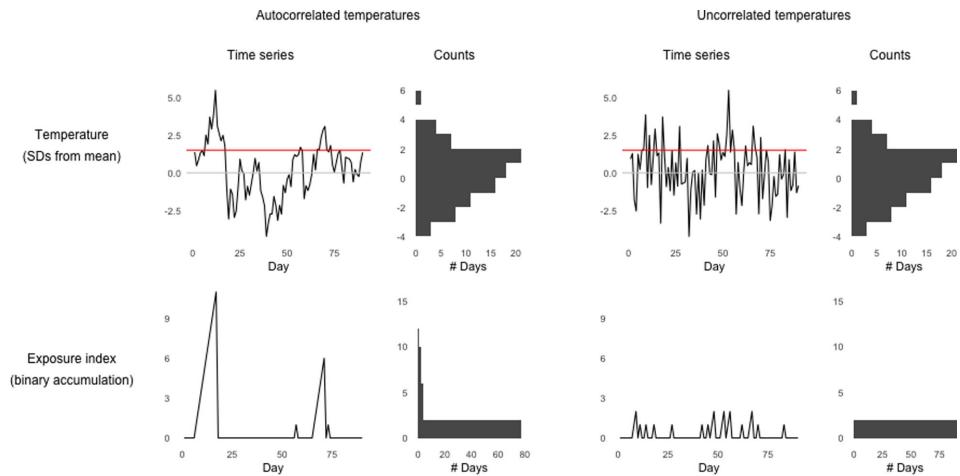


FIGURE 3. Illustration of heat accumulation models. Illustration of heat accumulation models as derived from temperature exposure. Line plots depict time series of temperatures (black) in standard deviations from the mean (grey), along with a +1.5 standard deviation abnormal heat threshold (red) (top plots), and the associated exposure index E_{lyd} through time (bottom plots). Histograms reflect counts of days for which temperatures (top) or the prolonged heat-exposure index (bottom) fall in different ranges, using the time series immediately to the left of that histogram. The left group of plots reflect an autocorrelated temperature series in which a heat wave occurs. The right set of plots use the same set of temperatures but in randomized order, so that no autocorrelation occurs.

a different time series of exposure indices for each location. The effect of the wave is simply a linear combination of the coefficients weighted by the changes in the number of days in each prolonged heat-exposure bin. We use this approach to understand the effects of three historical heat waves (Section 6.1.1) as well as projected damages under climate change (Section 6.1.2).

We illustrate the intuition and benefits of this approach in Figure 3 using our binary accumulation model and two hypothetical temperature series. The standard practice of counting days in temperature bins (top row) does not distinguish between a serially correlated temperature series with a heat wave (left group of plots) and the same set of temperatures in a re-randomized order (right group of plots). In contrast, both the time series and bin-based summary of our exposure index (bottom row) successfully differentiate between the two temperature series. The exposure index rises steadily when temperatures remain above T_{ld} (red line) for a prolonged period of time, giving rise to different distributions of E_{lyd} used to define our metric.

With this broad measurement framework in place, we return to the key threshold T_{ld} . We would like T_{ld} to identify days that are both abnormally hot for a location and time of year and hot in some intuitive sense. To this end, we require the temperature to be at least 1.5 standard deviations above average for a location and day and for average

historical temperatures on that day to be among the 90 hottest for that location.⁸ Using location-specific thresholds acknowledges that long-run adaptation to average climatic conditions is likely to influence the effects of a given temperature. Similarly, permitting the threshold to vary by day allows for seasonal acclimatization to hotter weather. We focus on the 90 hottest days because an abnormally warm day may have very different effects on output if the temperature is 50F (e.g. during winter) than if it is 100F (e.g. during summer), even if both are at least 1.5 standard deviations above average. Consistent with this approach, our intensity-based accumulation models use temperatures measured in standard deviations, so that the accumulation index E_{lyd} has units of (temperature) standard deviation days (SD-days). We also examine a range of alternatives for robustness (Online Appendix Table A.3), with the most important comparison to absolute temperature thresholds illustrating the importance of location specificity.⁹

To summarize, measuring prolonged exposure to heat in our framework requires three choices: a model of heat-stress dynamics, a summary of prolonged heat exposure across a period and region of interest, and a threshold for what constitutes an abnormally hot day. We emphasize results based on (1) the binary accumulation model, (2) the bin-based summary of exposure defined in equations (4) and (5), and (3) thresholds of 1.5 standard deviations above average for a location and day.

Before turning to our empirical application, we briefly note how our approach connects to prominent existing measures of heat waves or accumulated heat exposure. First, some prior work counts heat waves, defined as stretches of abnormally hot days longer than some minimum duration. Such a count is equal to the number of days with E_{lyd} exactly equal to that minimum duration under our binary accumulation model. For example, during each heat wave that lasts at least a week, E_{lyd} will equal 7 exactly once: on the 7th day of the wave. Second, a count of growing-degree days or cooling-degree days arises from a variant of our intensity-based model with $h_+(T_{lyd}, \bar{T}_{ld}) = \max(0, T_{lyd} - \bar{T}_{ld})$, $h_-(T_{lyd}, \bar{T}_{ld}) = 0$, and exposure summarized by the maximum value of E_{lyd} during the year. Third, the variant of our approach in which E_{lyd} is a length- D sequence of bin identifiers generalizes the standard approach of counting days in temperature bins (See Sections 2.2 and 6.3). We discuss other measures and their relation to our approach in Online Appendix Section A.10.

8. Specifically, we adopt a piecewise form for \bar{T}_{ld} :

$$\bar{T}_{ld} = \begin{cases} \mu_{ld}^T + 1.5\sigma_{ld}^T & \text{if } \text{rank}(\mu_{ld}^T) \leq 90 \\ \infty & \text{otherwise} \end{cases} \quad (6)$$

$$\text{where } \text{rank}(\mu_{ld}^T) = \#\{\mu_{ld'}^T : \mu_{ld'}^T \geq \mu_{ld}^T\}.$$

This threshold is in the spirit of some prior work, but our innovation is to compute averages and standard deviations using a window from 15 days before to 15 days after focal day d using data across all years in the sample. Massetti and Mendelsohn (2015) define climatology based on calendar months.

9. The distinction between location-specific and absolute (spatially-invariant) thresholds should be less important for analyses with smaller spatial extent.

3. Data

The focal dataset for this study is a global panel of daily temperature maxima from 1979–2016, gridded to 0.5 degree latitude and longitude cells. The data, compiled by the Climate Prediction Center (CPC), are based on raw observations from the Global Telecommunications System and are interpolated to the grid using the Shepard algorithm. Gridding of the data uses anomalies from monthly climatology; the latter is added back after to derive the final observations. For our main results, we exclude countries in Sub-Saharan Africa due to concerns over the spatial coverage, reporting frequency, and reliability of the stations on which the CPC dataset is based (see Online Appendix for details). Including those countries does not change our qualitative results (Table 4), but we have greater confidence in the filtered dataset.

For the second part of the paper, in which we examine relationships between heat waves and measures of economic output, we combine this temperature data with the data in Burke, Hsiang, and Miguel (2015). The primary control from Burke, Hsiang, and Miguel (2015) is annual precipitation (Matsuura and Willmott 2012). We also adjust for average annual temperature as in Burke, Hsiang, and Miguel (2015), but for internal consistency we derive average annual temperature from the CPC dataset rather than from Matsuura and Willmott (2012). The outcome of interest is growth in per-capita gross domestic product (GDP) (in constant 2010 US dollars per person) as provided in the World Bank's World Development Indicators. When examining differential impacts on agriculture and non-agricultural sectors, we use growth in value added per capita in those sectors, again from the World Bank. Studying output statistics of course has limitations; for example, heat waves may affect more than market-based activity, such as through the destruction of natural capital, which typically falls outside of national accounts. Further, some coping behaviors, for example, increasing electricity usage to power air conditioners, will show up as positive contributions to output, while the net welfare effects of those choices are unclear (e.g. due to increased pollution from electricity generation). Still, output serves as a useful outcome of interest for comparison with prior work.

To partly address the limitations of output statistics, we also introduce growth in the Food and Agriculture Organization (FAO) Crop Production Index as an alternate outcome. Examining a quantity-focused index allows us to examine the impacts of heat waves more completely. For example, if a major heat wave leads to large crop losses, price increases may follow, dampening effects of the heat wave on value added. While net effects on value are important for producers, reduced quantities may pose food security concerns.

As discussed in the preceding section, we summarize heat wave exposure at a country level by spatially averaging the per-location histograms of our exposure index. We do so in two ways, depending on the economic outcome of interest. When examining impacts of weather on overall or non-agricultural output growth, we weight exposure by year 2000 population density per grid cell from version 4 of the Gridded Population of the World dataset (CIESIN 2016). When examining growth in agricultural value added or quantities, we aggregate based on the proportion of each

cell devoted to crops or pasture in that same year. Fractions of 5 arcminute cells devoted to crops or pasture (Ramankutty et al. 2008) were obtained from EARTHSTAT and aggregated to the half-degree grid at which weather data are available. Precipitation and other temperature-based metrics use the same spatial aggregation scheme.

4. The Incidence of Heat Waves

Before studying the effects of heat waves on economic growth, we first illustrate how our heat-wave measures contain new information not present in previous measures of averages or counts of days in “hot” temperature bins. We do this by describing certain patterns of heat wave occurrence using our metrics and compare them with existing metrics.

We first describe some basic trends in heat-wave incidence across the globe (Figure 1). First, the incidence of days with high levels of prolonged heat exposure is increasing through time. Moreover, the incidence of higher exposure levels (9+, 12+) has increased at a faster rate. Because we define the threshold for what constitutes a hot day to be time-invariant for a location, this increasing incidence may reflect increasing variability, gradual warming that shifts the distribution of temperatures higher, or more clustering of hot days in time. However, the relative incidence of days between 90 and 100F and over 100F has not increased nearly as rapidly as that for heat waves, suggesting the growth in heat waves is not due to overall warming alone.

Despite these increasing trends, we note that long heat waves are relatively rare events, with countries experiencing much less than one such heat wave a year (Table 1). That low number reflects a combination of rarity and spatial extent: waves may often affect small parts of countries and country-scale waves may occur infrequently. The limited variation in heat wave exposure motivates our use of parsimonious specifications: we use a three-bin measure of heat-wave exposure in most models. Table 1 also highlights the important role of our relative thresholds: using fixed 90F or 100F thresholds identifies far more common events, with some countries surpassing those thresholds for the entire 90-day window on which we focus. When these rare heat waves do occur, Figure 3 also suggests that the El Niño Southern Oscillation (ENSO) may play a role. Average incidence is dramatically higher in the same year as or year after devastating El Niño events (e.g. 1982–83, 1997–98, 2002–03, 2009–2010, and 2014–2016). We discuss the potential value of this link in Section 7 and examine it in the Online Appendix (Section A.11).

An example baseline measure using our binary accumulation model corroborates geographic patterns of well-known heat-wave events (Figure 4). In 2003, Western Europe experienced a deadly heat wave, especially in France, while 2010 entailed a massive heat wave in Russia. The darker regions in Figure 4 show that the example heat-wave metric (days with $E \geq 9$ using binary accumulation) isolates these events reasonably well. These patterns of exposure to heat waves would not necessarily

TABLE 1. Heat wave summary statistics.

Duration	Crop/pasture-weighted				Population-weighted			
	Mean	Median	SD	Max	Mean	Median	SD	Max
Heat waves: 1.5 SD threshold								
Days $E \geq 6$	0.21	0.00	0.71	10.09	0.20	0.00	0.80	19.17
Days $E \geq 9$	0.05	0.00	0.29	7.50	0.05	0.00	0.38	11.38
Days $E \geq 12$	0.01	0.00	0.15	5.65	0.01	0.00	0.24	8.98
Heat waves: 90F threshold								
Days $E \geq 6$	16.90	6.48	22.65	85.00	14.21	4.05	21.41	85.00
Days $E \geq 9$	13.86	3.22	20.92	82.00	11.54	1.72	19.71	82.00
Days $E \geq 12$	11.86	1.62	19.49	79.00	9.80	0.79	18.33	79.00
Heat waves: 100F threshold								
Days $E \geq 6$	4.28	0.00	13.44	85.00	3.63	0.00	13.43	85.00
Days $E \geq 9$	3.56	0.00	12.16	82.00	3.07	0.00	12.27	82.00
Days $E \geq 12$	3.08	0.00	11.15	79.00	2.70	0.00	11.32	79.00
Abnormal heat thresholds (cell-level, unweighted)								
Mean + 1.5 SD (F)	86	89	18	126				

Notes: Summary statistics for spatially averaged counts of days falling in the specified prolonged exposure bin during the historically hottest 90 days for a particular cell. Rows correspond to different minimum prolonged exposure levels. Columns provide mean, standard deviation, and maximum. Row groups use different thresholds defining abnormally hot days: our primary threshold of 1.5 standard deviations above average for a cell and day of the year, as well as 90F and 100F for comparison.

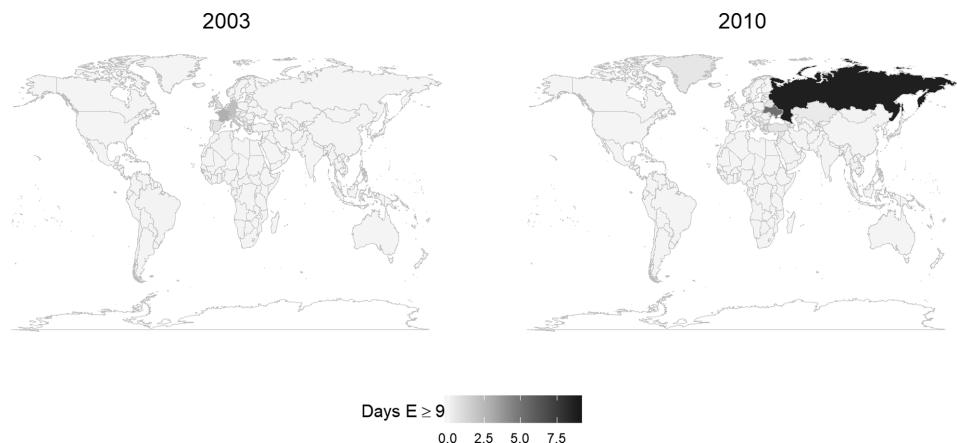


FIGURE 4. Metric performance for well-known heat-wave events. Measures of prolonged heat exposure in 2003 and 2010, during which major heat waves occurred in Western Europe (2003) and Russia (2010). Numbers represent crop- and pasture-weighted average number of days in which an individual in a country experienced prolonged heat exposure of at least nine (using the binary heat accumulation model). Darker regions represent greater prolonged exposure to heat. Accumulation uses a hot-day threshold of 1.5 standard deviations above normal.

be revealed by measures of temperature exposure used more commonly in applied econometric work (Online Appendix Figure A.6). More generally, the spatial patterns of heat-wave incidence using our metrics suggest a greater incidence of heat waves in countries that are wealthier¹⁰ and at temperate latitudes.

We illustrate the new information in our metrics using the same example as above: population density-averaged counts of days with $E \geq 9$. Correlations between this metric and both mean temperature and counts of days in hotter temperature bins (80–90F, 90–100F, >100F) are low. This is true whether we look at overall correlations (−0.04 to −0.01) or within-country correlations after removing static cross-country differences (−0.04 to 0.13). The latter is more relevant for the empirical application later, which employs per-country fixed effects. This suggests our approach captures new information not present in measures of heat exposure used in many economic analyses. In turn, the effects of heat waves may not be adequately captured by existing metrics.

The differences between our metric and counts of days in fixed temperature bins stem in part from the day- and location-specific nature of our thresholds. Our primary threshold defining abnormal heat (1.5 SD above average) has a cell-level median (mean) of 89 (86) degrees Fahrenheit—not far from a typical reference point of 90F. However, the standard deviation of our threshold is roughly 18F, and the interquartile range is 76F–98F. There is substantial variation across both space and time in what is considered abnormal, which is ignored in approaches using fixed temperature thresholds.

Given that our framework is general enough to admit a variety of metrics, we briefly summarize how several specific choices relate to one another (Online Appendix Table A.1). In short, correlations among metrics are affected strongly by the severity of the hot-day threshold (1, 1.3, 1.5, or 1.8 standard deviations above average) and by cutoffs for minimum prolonged exposure. Thus, two defining features of heat waves, abnormal heat and prolonged exposure, each play an important role in our constructed metrics. The contributions of both hot-day thresholds and sequencing are clarified by comparing heat-wave metrics with simple counts of hot days (ignoring sequencing) using identical hot-day thresholds. Counts of days with prolonged heat-exposure index of at least 3 are highly correlated (0.88) with the simple hot day counts, but days with exposure of at least 12 days exhibit much weaker correlation (0.27) (see Online Appendix Table A.2). Thus, the required sequencing of hot days imparts new information in those metrics. In contrast, using different samples (moving windows vs. calendar months) to determine thresholds or using different spatial aggregation schemes has more minor effects, as evidenced by high correlations across the metrics. The use of season restrictions (e.g. waves must occur during the historically hottest m days) has moderate effects on the final metrics from 90 to 150 days.

10. Wealthier countries are those with above-median 1980 GDP per capita (purchasing power parity).

5. Effects of Heat Waves on Economic Output

5.1. Empirical Framework

Building on an emergent climate–economy literature analyzing the impact of rising temperatures and extreme weather events (Dell, Jones, and Olken 2014), we estimate the effect of heat waves on growth rates of real GDP per capita (DY) using the following econometric model:

$$\begin{aligned} DY_{cy} = & \beta' W_{cy} + \alpha_1 \bar{T}_{cy} + \alpha_2 \bar{T}_{cy}^2 + \alpha_3 P_{cy} \\ & + \alpha_4 P_{cy}^2 + \mu_c + \eta_y + \eta_{c1}y + \eta_{c2}y^2 + \varepsilon_{cy}. \end{aligned} \quad (7)$$

Here W_{cy}] is a (potentially vector-valued) measure of heat-wave incidence in country c during year y , and the parameter vector β measures the effects of marginal increases in each element of the heat-wave vector W_{cy} .¹¹ When W_{cy}] is a vector of counts of days for which our prolonged heat-exposure index E_{lyd}] lies in different bins, each element captures the effect of an additional day with that exposure index compared to one with an exposure index in an omitted reference bin. Estimating the effect of an entire heat wave entails a linear combination of the elements of β that depends on the wave characteristics. In Section 6.1.1, we provide several examples of the impacts of heat waves in specific countries and years.

Our baseline specification conditions on the country’s annual mean temperature (\bar{T}_{cy}) and total rainfall (P_{cy}]), which are likely correlated with the occurrence of heat-wave events within the year. The inclusion of squared environmental variables captures potential nonlinearities in the relationship between temperature and growth (Burke, Hsiang, and Miguel 2015). Conditioning on mean temperature implies that the estimated heat-wave impact is net of idiosyncratic fluctuations in local average annual temperatures. While the two are related, a warmer year overall need not be tied to the occurrence of heat waves if extreme temperatures are not clustered consecutively or if they fall within relative (e.g. day- and area-specific) degree thresholds. Similarly, higher incidence of heat waves need not result in higher overall mean temperatures if temperatures are cooler in other parts of the year. Controlling for the mean temperature ensures that the estimated effect of the heat wave arises from the prolonged exposure of a minimum duration to warm days and not from a general warming or cooling overall.

In addition, the baseline includes country fixed effects (μ_c) and year fixed effects (η_y) whereby the former accounts for time-invariant country-specific characteristics that influence its rate of growth such as history, culture, and geography, while the latter captures year-specific worldwide shocks to output growth. We also allow for each country to exhibit its own level and quadratic trends in growth by the presence of linear and quadratic country time trends ($\eta_{c1}y + \eta_{c2}y^2$). They permit country growth rates to evolve non-linearly over time due to underlying features of the country’s economy, such as demographic transitions, institutional changes, and long-run income

11. We also allow for lagged heat wave effects in Section 6.3.3.

convergence. Using both country fixed effects and country-specific trends allow the trajectory of income levels in each country to exhibit a country-specific intercept, slope, and curvature (Hsiang and Jina 2014).

Given this framework, we identify the impact of heatwaves on output using within-country deviations from output trends. Identification rests on the frequency, duration, and intensity of heatwaves being exogenous conditional on our observed variables, fixed effects, and country-specific trends. Additionally, our model explicitly omits several covariates that determine growth, such as demographic or political variables, as they are themselves likely affected by climatic events and hence “bad controls” (Angrist and Pischke 2008). For inference, we allow the error terms (ε_{cy}) to exhibit heteroskedasticity and autocorrelation of unspecified form. The former may arise from differences in measurement precision by country while the latter may stem from any residual correlation in the outcome across time even after first differencing. We account for this by clustering standard errors at the country level.

Finally, while equation (7) reflects our baseline specification, we later employ minor variants to compare with prior work (Section 6.2) and to allow for longer-run effects and heterogeneity (Section 6.3). For the latter, we allow the impacts of heat waves to depend upon (a proxy for) development, whether a heat wave occurred earlier or later in the sample, and the average temperature in a year.

5.2. Differences in Growth When Heat Waves Occur

Examining simple differences in growth rates provides suggestive evidence that heat waves may impact growth, particularly in the agricultural sector (Table 2). We first demean annual growth rates within each country, then compare differences in those demeaned growth rates in countries and years with any heat waves and those without. Observations are included in the heat-wave group if at least one location in that country experiences at least one day during the year at or above the specified minimum exposure using the binary accumulation model. The (demeaned) average agricultural GDP growth rates in countries experiencing heat waves are significantly lower for two of three wave definitions, while there is little compelling evidence of a comparable difference in other sectors. To more carefully examine whether these differences are indeed due to waves, we turn to our regression results.

6. Results

We first offer a brief roadmap to the results to follow. We begin by illustrating the effect of an additional day of prolonged heat exposure on output (Section 6.1). We then ask whether our results hold under alternate sets of assumptions, estimating models that vary the controls used, the sample, and our “abnormal heat” threshold. After establishing the robustness of our results, we quantify damages from heat waves in two ways. First, we use historical and constructed counterfactual temperature profiles to estimate the damages from three high-profile heat waves that occur in our dataset (Section 6.1.1). Second, we use a suite of 21 climate models offering daily projections to examine potential impacts of heat waves from 2020–2100 (Section 6.1.2).

TABLE 2. Average demeaned growth rates with and without heat waves.

Min. Exposure	No waves			Waves			Diff.
	Mean	(SD)	N	Mean	(SD)	N	
Overall							
$E \geq 6$	0.0002	(0.0645)	2,407	-0.0003	(0.0542)	1,764	-0.0005
$E \geq 9$	0.0005	(0.0628)	3,394	-0.0021	(0.0483)	777	-0.0026
$E \geq 12$	0.0003	(0.0617)	3,861	-0.0036	(0.0413)	310	-0.0039
Non-Agriculture							
$E \geq 6$	-0.0007	(0.0558)	1,897	0.0009	(0.0533)	1,536	0.0016
$E \geq 9$	-0.0001	(0.0572)	2,743	0.0003	(0.0434)	690	0.0003
$E \geq 12$	0.0002	(0.0557)	3,160	-0.0028	(0.0414)	273	-0.0031
Agriculture							
$E \geq 6$	0.0033	(0.0818)	1,808	-0.0041	(0.0881)	1,469	-0.0075**
$E \geq 9$	0.0021	(0.0845)	2,636	-0.0087	(0.0852)	641	-0.0109**
$E \geq 12$	0.0011	(0.0848)	3,027	-0.0132	(0.0829)	250	-0.0143**

Notes: Within-country demeaned overall, non-agriculture, and agriculture growth rates for country-year observations with no heat waves versus those with heat waves, as well as differences in demeaned growth rates between those two groups. Row groups correspond to different types of growth, and individual rows reflect different minimum prolonged exposure levels (using the binary accumulation model) that must be reached to count as a heat wave. Stars indicate significant differences in demeaned growth for countries with and without heat waves. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We then directly compare our findings to those which arise if using existing measures of heat exposure (Section 6.2). We do so first for the common practice of counting days which fall in fixed temperature bins, ignoring accumulation (Section 6.2.1). Second, we compare our findings to those from counting stretches of consecutive hot days, a more common practice outside of economics (Section 6.2.2).

Our final set of results (Section 6.3) examines extensions to our approach, studying whether prolonged exposure to heat impacts economies in ways that are not captured by the motivating model in Section 2.2. Specifically, we estimate models that allow for heat waves to have longer-run effects, to alter more than just TFP, and to have impacts that may vary with development, average temperature, and time. Finally, we examine whether sequences of heat exposure other than uninterrupted stretches are still damaging.

6.1. Main Findings: Heat Waves Reduce Output

We find strong evidence that heat waves occurring during the hottest parts of the year depress economic output in the agricultural sector, with more modest evidence of an impact in other areas of the economy (Figure 5).¹² Our most concerning estimates

12. In the Online Appendix, we reproduce our main results but use standard errors that allow for cross-country spatial correlation per Conley (1999) and within-country autocorrelation. Inference is largely unchanged; see Online Appendix Figure A.1 and the accompanying caption.

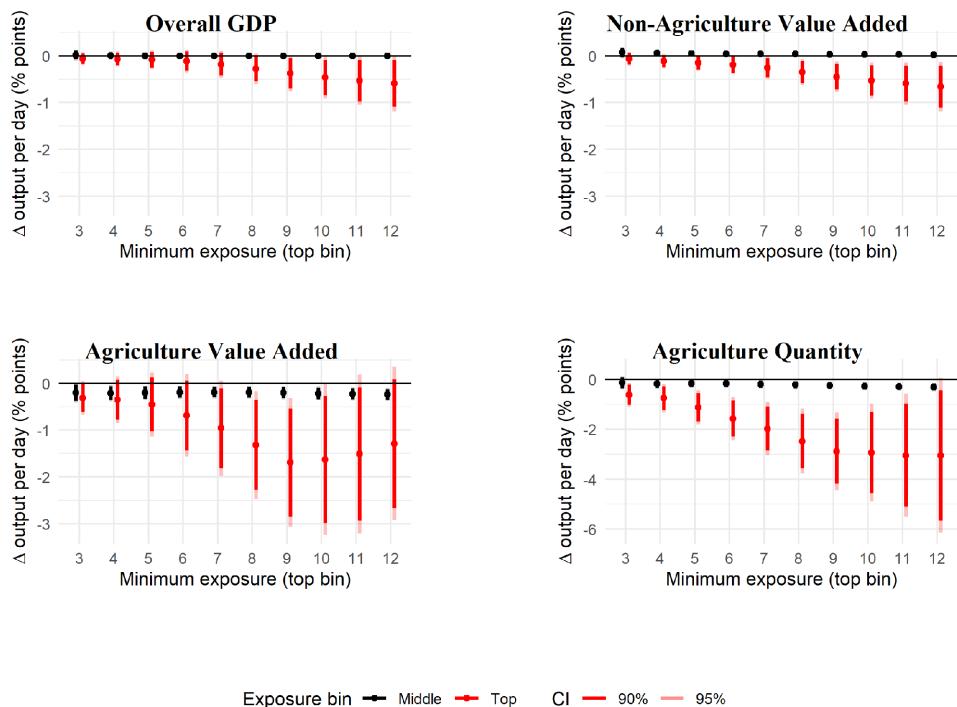


FIGURE 5. Effects of prolonged exposure to heat. Effects of an additional day during with prolonged heat-exposure index in the specified bin during the historically hottest 90 days of the year. Each location on the horizontal axis represents a different specification. Each specification uses three bins to measure prolonged heat exposure: an omitted bottom reference bin with $E = 0$, a middle bin with $1 \leq E < C$ for C equal to the location on the horizontal axis, and a top bin with $E \geq C$. All estimates use an abnormal heat threshold of ≥ 1.5 SD above the local average temperature during the historically hottest 90 days for a location. Panels depict different outcomes. 90% confidence intervals shown are based on standard errors clustered at the country level.

suggest that an abnormally hot day preceded by at least eight others reduces per-capita agricultural value added by approximately 1.7%. The effect of prolonged exposure to heat on a quantity measure of output in agriculture is stronger. Specifically, an abnormally hot day preceded by at least eight others reduces the FAO Crop Production Index by almost 3%. The stronger effect on quantities is consistent with supply shocks inducing price increases, such that the effect on value added is dampened. Results using intensity-based accumulation of heat follow similar patterns; we briefly review those results in Section 6.3.1.

These results reflect a heat wave effect conditional upon—that is, not captured by—mean temperature.¹³ As illustrated for our binary accumulation model (Table 3), our estimates of the effect of mean temperature and its square are also consistent with

13. Annual mean temperature is, strictly speaking, a “bad control.” We include mean temperature because (1) the effect of a heat wave on the annual mean should be small, and (2) we aim to isolate an effect that is

TABLE 3. Example effects of other covariates on growth.

Coefficient	Overall	Non-agriculture	Agriculture
Days $1 \leq E \leq 8$	0.00001 (0.00031)	0.00035 (0.00025)	-0.00198*** (0.00071)
Days $E \geq 9$	-0.00377* (0.00195)	-0.00448*** (0.00165)	-0.01691** (0.00702)
Mean Temp.	0.01190*** (0.00372)	0.00875 (0.00610)	0.02524*** (0.00677)
Mean Temp. ²	-0.00033*** (0.00011)	-0.00032* (0.00018)	-0.00078*** (0.00023)
Precip.	-0.00350 (0.01361)	-0.01819 (0.01207)	0.02981 (0.03254)
Precip. ²	-0.00084 (0.00282)	0.00330 (0.00247)	-0.01286 (0.00776)
N	4171	3433	3277

Notes: Estimated relationships between covariates and growth for a baseline model counting the number of days in which the prolonged heat-exposure index (E) falls in the ranges $[1,8]$ or $[9, \infty)$ using the binary accumulation model and a $+1.5$ SD threshold. All specifications include country fixed effects, year fixed effects, and quadratic country time trends, which are not reported. Temperature is measured in degrees Celsius and precipitation in meters. Numbers of observations differ due to availability of both sector-level outcomes and data underlying different spatial weighting schemes. Standard errors are clustered at the country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Burke, Hsiang, and Miguel (2015); that concordance corroborates our earlier claim, based on variable correlations alone, that our heat-wave measures introduce genuinely new information. In sum, we find evidence that both heat waves and warming appear to matter for economic growth; our results complement rather than challenge prior work.

The effects of prolonged heat exposure remain stable across a suite of robustness checks, summarized in Table 4. We first consider three alternate samples: (1) countries having more than 20 years of data to ensure adequate data for per-country trends, (2) including Sub-Saharan Africa despite our data concerns, and (3) countries for which the threshold defining abnormally hot days (averaged across days and grid cells in a country) is at least 80F. We also consider modified specifications, first using only linear per-country trends, then no such trends, including a control for lagged growth, and adjusting for numbers of days in different 10 degree (F) temperature bins rather than mean temperature and its square. These exercises present no evidence that contradicts our baseline findings. Agricultural growth remains negatively affected by heat waves, and we continue to find weaker evidence of an effect of heat waves on non-agricultural and overall GDP growth.

As a complement to these sample-based robustness checks, we also conduct falsification tests based on randomization (Online Appendix Figure A.3). We first randomize average temperature to examine whether our heat-wave estimates actually

distinct from that documented in Burke, Hsiang, and Miguel (2015), even if our partial effect differs from the total effect.

TABLE 4. Robustness to alternate samples and specifications.

	Baseline	>20 yr	With SSA	$T_{ld} \geq 80^{\circ}\text{F}$	Lin. trend	No trend	Lag growth	T bins ctrl
Overall GDP growth								
Days with $1 \leq E \leq 8$	0.0000	0.0001	0.0002	-0.0001	0.0000	0.0000	0.0000	0.0001
Days with $E \geq 9$	-0.0038*	-0.0043**	-0.0032	-0.0038*	-0.0024	-0.0022	-0.0041**	-0.0041*
Non-Agriculture GDP growth								
Days with $1 \leq E \leq 8$	0.0003	0.0003	0.0003	0.0003	0.0002	0.0004	0.0005*	0.0003
Days with $E \geq 9$	-0.0045***	-0.0044***	-0.0038**	-0.0047***	-0.0028	-0.0027	-0.0044**	-0.0054***
Agriculture GDP growth								
Days with $1 \leq E \leq 8$	-0.0020***	-0.0019**	-0.0020***	-0.0020***	-0.0020***	-0.0019***	-0.0025***	-0.0010
Days with $E \geq 9$	-0.0169**	-0.0170**	-0.0101	-0.0180**	-0.0165**	-0.0154**	-0.0123**	-0.0158**
# observations								
Overall	4171	3348	5757	3835	4171	4122	4171	4171
Non-Agriculture	3433	3083	4749	3154	3433	3365	3433	3433
Agriculture	3277	3083	4490	2890	3277	3212	3277	3277

Notes: Each estimate is from a separate regression. Except where noted, all specifications include country fixed effects, year fixed effects, and quadratic country time trends. All specifications include controls for mean temperature, temperature squared, precipitation and precipitation squared. Temperature is measured in degrees Celsius and precipitation in meters. Standard errors are clustered at the country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

reflect average temperature effects for which we inadequately control. Our estimates change little, suggesting that is not the case (panel a). We next randomize our heat-wave measures, which should yield null results, which it does (panel b). Repeating that exercise 1,000 times shows our actual estimates fall well outside the range of estimates that arise if the link between heat waves and output is broken (panel c). In sum, these randomization tests do not contradict our main findings, lending credibility to our results.

Finally, we also examine robustness to how we define an “abnormally hot” day in a location. The results of these exercises, summarized in Online Appendix Table A.3, reveal three sensible patterns.¹⁴ First, our use of location- and day-specific thresholds is critical: we see only limited effects of waves when defining abnormally hot days based on absolute thresholds (90F or 100F). Second, stricter requirements for a hot day, in the form of higher multiples of standard deviation or shorter season restrictions (e.g. hottest 90 vs. hottest 120 days), intuitively lead to stronger estimated effects of heat waves. Finally, our baseline results are robust to how “abnormal” is defined, whether we use months rather than moving windows or day-specific quantiles rather than standard deviation-based thresholds.

6.1.1. Estimated Lost Output: Historical Estimates. To put our estimates in context, we quantify lost output from three high-profile heat waves: 2003 in France, 2010 in Russia, and 2012 in the United States. In each case, we construct a counterfactual temperature history for the country and year in which abnormally hot days (≥ 1.5 SD above average) instead had average temperature. Comparing the actual and counterfactual temperature histories yields differences in our heat-wave measure and temperature measures, which we combine with marginal effects (Figure 5) to yield overall damage estimates.

Our primary heat-wave measure implies per-wave reductions in output ranging from \$0.8–3.1 billion for agriculture and up to \$31.9 billion in other sectors (Table 5).¹⁵ To help interpret these estimates, we narrow in on France in 2003. Actual agricultural output in France declined \$6.2 billion (15.2 percentage points) from 2002 to 2003, \$3.1 billion (7.6 percentage points) of which we attribute to the summer heat wave. In contrast, a model using absolute temperature bins (see Section 6.2) indicates losses of \$318 million (0.77 percentage points), while simple counts of hot days with no explicit treatment of heat waves implies a \$1.1 billion loss (2.7 percentage points). The contrast is equally stark when we estimate total lost output from the three heat waves we consider: our primary heat wave specification suggests \$52 billion, while a standard temperature-bin approach implies losses of only \$13.3 billion. Our

14. Online Appendix Table A.3 uses the binary accumulation model to facilitate comparisons with absolute and quantile temperature thresholds. For those thresholds, exceedance of the threshold would be naturally measured in degrees rather than standard deviations used in our primary 1.5-sd threshold. The binary accumulation model offers a more consistent comparison.

15. For reference, the National Oceanic and Atmospheric Administration puts overall damages from the 2012 heat wave in the US at \$34.2 billion (<https://www.ncdc.noaa.gov/billions/events>). Our binary accumulation heat wave specification suggests a loss of \$0.8 billion in agriculture and an impact in other sectors that is not significantly different from zero.

TABLE 5. Estimated lost output from heat waves (\$B).

Heat Wave	Temp. Bins		# Hot days		Prolonged heat exposure	
	Ag	Non-Ag	Ag	Non-Ag	Ag	Non-Ag
France 2003	-0.3 (-0.4, -0.2)	-0.5 (-5, 4)	-1.1 (-1.6, -0.6)	3.3 (-6.9, 13.6)	-3.1 (-5.2, -1)	-31.9 (-58.6, -5.2)
Russia 2010	-0.8 (-1.1, -0.4)	0.1 (-5.5, 5.7)	-3.1 (-4.4, -1.7)	4.2 (-8.8, 17.3)	-3.6 (-6, -1.2)	-18 (-34.3, -1.6)
USA 2012	-0.4 (-0.7, -0.1)	-11.4 (-28.7, 5.9)	-4.8 (-6.9, -2.6)	22.5 (-46.9, 92)	-0.8 (-1.3, -0.3)	5 (-11.7, 21.7)

Notes: Estimated reductions in output (in billions of 2010 USD) from heat waves for 3 years in which high-profile heat waves occurred. Numbers in parentheses represent 90% confidence intervals for the estimate above. The first two result columns come from a model using a standard count of days in 10-degree temperature bins, while the third and fourth columns come from a hybrid approach counting days at least 1.5 standard deviations above the long-run average for a location and day of the year. The fifth and sixth result columns reflect our measure of prolonged heat exposure based on binary accumulation of heat, wherein a day that is not abnormally hot provides complete relief from prior heat exposure ($\gamma = \infty$).

approach thus suggests damages from heat waves may be several times larger than might otherwise be suggested by more conventional estimation techniques. We more formally compare these approaches in Section 6.2.

6.1.2. Estimated Lost Output: Climate Model Projections. As climate change is likely to alter the incidence of heat waves (Meehl and Tebaldi 2004; Cowan et al. 2014), we also pair our results with a set of projections from climate models to quantify potential future impacts from heat waves. We focus our attention on future agricultural impacts for two reasons. First, that sector is more heavily affected by heat waves. Second, we find little evidence that agricultural impacts of heat waves have changed during our sample (Section 6.3.2 and Online Appendix Table A.8), lending some credibility to the projection exercise. Even so, because further adaptation could cause future agricultural impacts to diverge from those in the past (see Section 6.3.2 for a discussion), we compute projections under two contrasting assumptions about adaptation, as explained below.

To assess future heat-wave impacts, we use 21 models from the Coupled Model Intercomparison Project Phase 5 for which daily, downscaled temperature predictions have been made available in the NASA Earth Exchange Global Daily Downscaled Projections dataset (Thrasher et al. 2012). For each model, we compute our heat-wave metrics using the binary accumulation model from 2020 to 2100 for RCP 4.5.¹⁶ To quantify the change in a country's heat-wave incidence due to climate change, we compare the modeled incidence in a year to average incidence in that country in the first 5 years of the projections. Pairing that difference with corresponding estimates

16. RCP8.5 produces heat-wave incidence that is far outside what is observed in our data, exacerbating concerns over the linear and reduced-form nature of our econometric model. For example, RCP8.5 can yield more than 100% losses for some climate models. For these reasons, we focus on RCP4.5.

underlying Figure 5 yields our estimated climate impacts.¹⁷ For comparison, we undertake analogous steps to estimate the effects of mean temperature shifts.

We conduct this exercise under two contrasting assumptions about future adaptation. First, we assume no adaptation, using the same estimates and thresholds defining what constitutes abnormal heat for a location and time of year. This assumption is consistent with the absence of significant evidence of changing impacts of heat waves on agriculture thus far (Section 6.3.2 and Online Appendix Table A.8),¹⁸ but may not adequately reflect future adaptation efforts. To that end, our second set of projections assumes that the impact of a stretch of abnormally hot days will remain unchanged, but that thresholds defining abnormal heat will change gradually in the future. Specifically, for each location and day of the year, we allow the mean and standard deviation of temperature to evolve at the annual rates they have in historical data. In many locations, this assumption means days will have to be increasingly hot in order to contribute to a heat wave in the future. This is not the only way to model adaptation, but it is a logical extension of our heat wave definition.

Results of these projections are depicted in Figure 6. Using model ensemble mean projections, average per-country losses reach 10.3% of agricultural output per year by 2091–2100 without adaptation, and 4.5% with adaptation. Thus, should future adaptation alter thresholds for abnormal heat in the manner described above, adaptation could reduce end-of-century heat wave damages by approximately 56%, though potential losses remain concerning. In contrast, the mean effect of warming by that time is only 0.9%, with many initially cooler countries (e.g. Russia) even experiencing small gains. Heterogeneity is also evident in the projected impacts of heat waves, both in the country-specific profiles and in the long tail in ensemble prediction density. In short, the impact of climate change on agriculture through heat waves is potentially quite large and likely to vary substantially across countries.

As with any projections of climate damage, our estimates have limitations. In addition to the challenges inherent in the climate models on which we rely, the extent and efficacy of future adaptation are difficult to predict. The increasing incidence of heat waves (Figure 1) should raise expected returns to adaptation. However, as discussed in Section 7, adapting to rare events like heat waves poses distinct challenges.

6.2. Comparison With Other Approaches

We next contrast our approach with other ways to study the effects of temperature. As discussed in Section 6.1 and depicted in Table 3, our findings complement prior evidence of a relationship between mean temperatures and output. Here we turn to other common practices for studying the effects of temperature: counting days that

17. As described in Section 6.3.3, we find only evidence of level effects, not growth effects. As such we do not incorporate other growth projections in these estimates.

18. Limited effectiveness of past agricultural adaptation to heat has also been noted elsewhere (e.g. Auffhammer and Schlenker 2014; Burke and Emerick 2016).

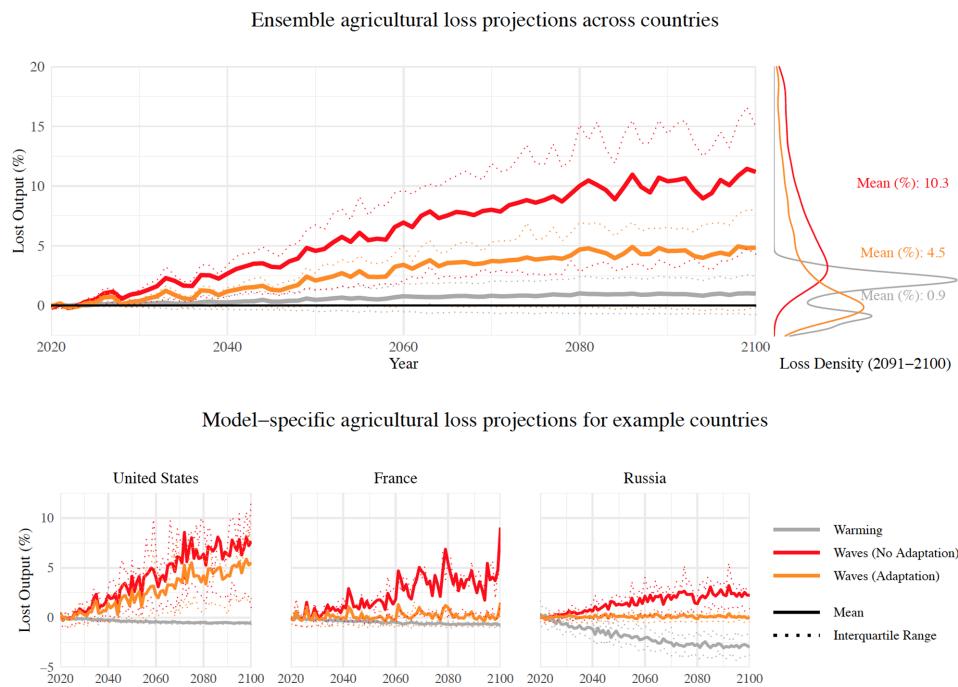


FIGURE 6. Projected agricultural output losses from heat waves. Point estimates for agricultural output losses (in %) per year. Colors indicate damages for heat waves under a fixed threshold for abnormal heat (red) and one that shifts based on historical trends (orange), as well as damages due to warming (gray). Shifting thresholds for abnormal heat approximate changes in damages if society adapts to changing temperature regimes. Top panel: ensemble mean losses as a function of time, with solid lines depicting global means and dotted lines depicting the interquartile range of ensemble mean losses across countries. Cross-country distributions and means of losses in the final decade (2091–2100) are depicted to the right. Bottom panels: ensemble means (solid lines) and cross-model interquartile range (dotted lines) of projected agricultural losses for each country examined in Table 5.

fall in temperature bins and counting continuous stretches of hot temperatures longer than a minimum duration.

6.2.1. Counting Days in Temperature Bins. Our main estimates imply heat-wave effects that are much larger than would be obtained using conventional approaches based on temperature bins. We base this claim on two types of evidence. First, we estimate models that count days in discrete temperature bins without regard for when those days occur.¹⁹ One model employs commonly used ten-degree Fahrenheit bins, with 60–70F as an omitted reference bin. Because using fixed temperature bins ignores location- and season-specific climatology, our second model includes counts of days in one of two bins split by our “abnormally hot” temperature threshold of 1.5

19. These models exclude mean temperature and its square, since temperature bins are frequently used as an alternative to estimated moments of temperature distributions.

TABLE 6. Effects of hot days, ignoring sequencing.

Bin	Overall	Non-Agriculture	Agriculture
Conventional Temperature Bins			
70–80	0.0001 (0.0002)	0.0003** (0.0001)	−0.0005 (0.0004)
80–90	0.0003 (0.0002)	0.0004** (0.0002)	−0.0011*** (0.0004)
>90	0.0001 (0.0002)	0.0003* (0.0002)	−0.0014*** (0.0004)
# Hot Days (Local and Seasonal 1.5 SD Threshold)			
# Hot Days	−0.0002 (0.0003)	0.0001 (0.0003)	−0.0026*** (0.0007)

Notes: Two sets of models examining the role of hot days separately from the effect of heat waves. Top: effects of a day in the specified temperature bin relative to a day in the 60–70F bin. Colder bins are omitted for brevity; full results are reported in Online Appendix Table A.12. Bottom: effects of a day above our location- and day-specific temperature threshold of 1.5 SD above average. Numbers of observations by outcome are as in Table 3. Standard errors are clustered at the country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

standard deviations above average. Still, both approaches ignore the time structure of exposure to heat, focusing only on the number of days considered hot by either absolute or locally-defined thresholds.

Counting hot days without regard for when they occur suggests that heat waves have much smaller effects than our main estimates indicate (Table 6). The top group of results suggest limited negative effects of an additional day in 80–90F or >90F temperature bins on agricultural output, with much smaller and even positive effects of those hotter days on other sectors (colder bins are omitted for brevity but are available in Online Appendix Table A.12). Each day above 90F reduces agricultural output by 0.14 percentage points compared to a day of 60–70F weather. Even after defining heat thresholds based on local climate, the bottom set of results suggest that a day more than 1.5 standard deviations above average reduces agricultural GDP growth by 0.26 percentage points. Both results contrast sharply with the estimates from our main models that account for heat accumulation (Figure 5). For example, using the binary accumulation model, each day with prolonged heat-exposure index at or above nine reduces agricultural output by roughly 1.7%—several times larger than the biggest effect reported in Table 6.

To more formally test whether prolonged exposure to heat matters, we use our binary accumulation model to ask whether the effect of an abnormally hot day depends on whether that day is preceded by other abnormally hot days. Specifically, we divide prolonged heat exposure into four bins: (0, 1), (1, 2), (2, 9), and (9, 90). By defining the (1, 2) bin to be the omitted reference bin, all estimates in the model can be interpreted as the effect of an additional day in the specified prolonged heat exposure bin as compared to a day in the (1, 2) bin. For the chosen accumulation model, a day in the (1, 2) reference bin is an abnormally hot day not preceded by other hot days, so that coefficients associated with other bins indicate whether there are non-additive

TABLE 7. Differential effects of hot days preceded by other hot days.

Exposure level	Overall	Non-agriculture	Agriculture
$E = 0$ hot days	0.0003 (0.0009)	0.0001 (0.0010)	0.0031* (0.0016)
$E = 2\text{--}8$ hot days	0.0006 (0.0015)	0.0008 (0.0016)	0.0018 (0.0029)
$E \geq 9$ hot days	-0.0038* (0.0021)	-0.0049*** (0.0017)	-0.0151** (0.0072)

Notes: Effect of a day in the specified prolonged heat exposure bin on output growth overall (column 1) or in particular sectors (columns 2 and 3). The omitted reference bin represents a day in the (1, 2) bin which, for the binary accumulation model, is an abnormally hot day preceded by at least one day below the 1.5 SD threshold. Numbers of observations by outcome are as in Table 3. Standard errors are clustered at the country level.
 * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

effects of prolonged exposure to heat. The top bin of (9, 366) is chosen for continuity in our running interpretation of a 9-day wave.

The results of this exercise indicate that abnormally hot days toward the end of long heat waves have a significantly different effect than those at the start of such waves (Table 7). Specifically, an abnormally hot day preceded by at least eight other such days (i.e. the ninth or greater day of a heat wave) reduces agricultural GDP growth by approximately 1.5 percentage points more than an abnormally hot day preceded by at least one cooler day.²⁰ Differences for other sectors of the economy remain significant, but are substantially smaller, consistent with earlier results. Finally, the positive estimates in the top row of Table 7 indicate that cooler days (in prolonged heat exposure bin (0, 1) are better for growth than isolated hot days, with magnitudes consistent with the analogous comparison in the final row of Table 6.

Altogether, we interpret this evidence to mean that abnormally hot days are bad for growth—particularly agricultural growth—but hot days preceded by many other hot days are far worse. Further, the relationship between prolonged exposure E and its effect on growth is nonlinear: a few hot days in a row ($E \in [2, 8]$) are not significantly better or worse for growth than an isolated hot day ($E = 1$). Only after much more prolonged exposure do effects on growth intensify. In short, heat waves matter.

6.2.2. Counting Heat Waves Longer Than a Minimum Duration. We next compare our approach, based on heat accumulation models, to the practice of counting waves of some minimum duration. Specifically, we estimate variants of our main specification that include both a count of waves at least D days long and the number of days with $E \geq D$ under our binary accumulation model. The coefficient on our prolonged heat-exposure index then reflects the damage of an additional day in a heat wave, holding the number of waves fixed.

20. Redefining the top bin as cumulative exposure of 8 (10) or more days yields similar results: an estimated reduction in agricultural growth by 1.1 (1.4) percentage points.

TABLE 8. Effects of wave duration beyond minimum thresholds.

	Overall	Non-agriculture	Agriculture
6+ day wave			
Per wave	0.0093 (0.0069)	0.0090 (0.0059)	0.0044 (0.0245)
Per day above threshold	−0.0034* (0.0018)	−0.0031* (0.0016)	−0.0132** (0.0062)
9+ day wave			
Per wave	0.0080 (0.0105)	0.0088 (0.0091)	−0.1221** (0.0570)
Per day above threshold	−0.0052* (0.0027)	−0.0049** (0.0024)	0.0013 (0.0117)

Notes: Tests of damage caused by additional days above a specified heat accumulation threshold as compared to simple counts of the number of heat waves. Numbers of observations by outcome are as in Table 3. Standard errors are clustered at the country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

This exercise illustrates that duration provides additional information relevant to output not captured in simple counts of heat waves, especially outside of agriculture (Table 8). As intuition would suggest, all statistically significant effects of increasing a wave's duration are negative. These results are also consistent with the findings in Table 7, but allow for more direct comparison with the simple approach of counting waves.

6.3. Extensions

6.3.1. Intensity-Based Accumulation. Our intensity-based accumulation model, which incorporates the severity of heat alongside the duration of a heat wave, yields qualitatively similar results to the binary accumulation model we have emphasized thus far (Online Appendix Figure A.2). Output declines when a year contains periods of heat that are longer or more intense, and losses are larger in agriculture than in other sectors of the economy. In the extensions we investigate in the following sections, our findings remain insensitive to the choice of accumulation model. We continue to focus our discussion on the binary accumulation results for brevity, but present results for the intensity-based model where relevant (Online Appendix, Tables A.4–A.10)

6.3.2. Heterogeneity. In addition to the average effects estimated above, we briefly examined three ways in which the impacts of heat waves vary through the use of interactions. Specifically, we examined whether the effects of heat waves depend upon (1) a country's level of development as proxied by start-of-sample GDP per capita, (2) the year in which a heat wave occurred (via a linear trend in impacts), and (3) whether the year in which a wave occurred was hotter than normal for the country in question. These investigations have several motivations. First, certain economies may be better equipped to handle heat waves than others, and our analytical framework permitted the relationship between prolonged heat exposure and TFP

(f_{il}) to be location-specific. Second, while heat waves are projected to become more common, adaptation over time could dampen their effects. Third, while our analytical framework used an assumption of additive separability to simplify the motivation for our approach (equation (1)), if that assumption does not hold, heat waves could have different effects in years that are warmer overall.

Across these dimensions we find evidence consistent with intuition, albeit of mixed strength. The strongest evidence of non-constant effects of heat waves arises outside of the agricultural sector: we find that waves are more damaging (1) in poorer countries, (2) in earlier years, and (3) in years that are hot compared to a country's long-run average temperature (Online Appendix, Tables A.6–A.8). The first two findings are consistent with but not necessarily causal evidence of adaptation. Individuals and firms in wealthier countries may have access to technology and capital to cope with waves (e.g. air conditioners), and investment in such capital may have increased in response to heat waves in earlier years. The smaller impact of more recent heat waves is also broadly consistent with evidence that the effects of heat waves on mortality in the United States have declined over time (Barreca et al. 2016).

We find little comparable evidence of heterogeneity in the agricultural sector. The few significant interactions echo our findings in other sectors that heat waves are more damaging in years that are warm overall. In contrast, we find no evidence that the impact of heat waves on agriculture has changed through time. The absence of a discernable trend should be interpreted with caution, as it could arise for several reasons. First, farmers may have chosen not to adapt if costs outweighed benefits after adjusting for assessment of probabilities and risk preferences. Second, while Figure 1 indicates increasing wave incidence and a potential link to ENSO, waves in particular locations and years remain difficult to predict well in advance. That could hinder formation of expectations and hence assessment of the benefits of adaptation. Third, farmers may have adapted, but those efforts may have been offset by what would otherwise be increasingly damaging heat waves. Fourth, our empirical approach, which examines linear trends in impact using fixed effects models, could miss some forms of adaptation. Adaptation could result in nonlinear changes in heat wave impacts, or some of the benefits of adaptation could be subsumed by year effects and per-country trends. Each of these possible channels leaves room for more successful future adaptation: costs could decline with new technologies, benefits could rise if waves become more damaging, or forecasts of heat waves could become more reliable at longer lead times. It is for these reasons that we conduct alternate climate projections in Section 6.1.2 in which criteria for what constitutes a heat wave changes through time.²¹

6.3.3. Longer-Run Effects. While our main estimates in Section 6.1 show clearly that heat waves can reduce output, they focus only on short-run, contemporaneous impacts. We also examine longer-run effects of heat waves using both finite distributed lag (FDL) and autoregressive distributed lag (ARDL) models. Both allow for heat

21. We thank an anonymous reviewer for suggesting this exercise.

waves and other environmental variables in past years to affect current growth, with the ARDL model also allowing for dynamics in output growth. Combinations of estimates from each model can be used to test whether heat waves have persistent impacts on growth or simply short-run level effects.

As indicated in Online Appendix Tables A.4 and A.5, we find little evidence that heat waves have significant longer-run impacts on growth using either ARDL or FDL models. Most estimated long-run effects are near zero and imprecisely estimated. This is true regardless of which output we consider and whether or not our heat-accumulation model depends on intensity. The lack of long-run effects is perhaps unsurprising, especially for agriculture. Farmers experiencing crop loss from a heat wave in one year can simply replant the next, causing a rebound in output and hence growth rates.²²

6.3.4. Effects of Other Sequences of Temperature Exposure. The accumulation models on which we focus assume a single cooler day provides complete relief from prolonged exposure to heat. While this assumption matches colloquial understanding of what constitutes a heat wave (consecutive hot days), we briefly extend our approach in two ways to study the impacts of other patterns of exposure to abnormal heat.

First, we note our framework is general enough to allow for cooler days to provide only partial relief from prolonged heat exposure. Specifically, the function $h_-(T_{lyd}, \bar{T}_{ld})$ may be smaller in magnitude than E_{lyd-1} , such that a cooler day provides only partial relief from recent heat. For example, we might define $h_-(T_{lyd}, \bar{T}_{ld}) = \gamma(\bar{T}_{ld} - T_{lyd})$, with γ scaling the relief provided by a cooler day. In this framework, our main estimates arise from a model in which $\gamma = \infty$. However, models in which cooler days provide only partial relief from prolonged heat exposure (e.g. $\gamma \in \{1, 3, 7\}$) are useful for at least two reasons. First, estimates from models with partial relief offer an informal falsification test for our main findings. When γ is smaller, our index E can reach a given level with hot and cold days interspersed, so that our estimates reflect events that are intuitively less severe. If we are indeed measuring the effects of temperature and not some omitted driver of output, we should expect effect sizes to be smaller in magnitude for a fixed value of E when γ is smaller. Second, if we are interested in prediction rather than inference, selecting a “best” value of γ that most closely reflects physical processes of heat stress may be of interest.

Estimates from models in which cooler days provide partial relief follow sensible patterns. When cooler days provide only partial relief, the effects of a day with E greater than a given threshold are indeed smaller in magnitude (Online Appendix Table A.9). More formal comparison reveals differences in estimated effects between $\gamma = 1$ and $\gamma = 366$ are statistically significant in agriculture (Online Appendix Table A.10). Turning to the question of which γ has the best predictive performance, both AIC and BIC select $\gamma = 366$ for agriculture but $\gamma = 1$ for other

22. This is no different from other dramatic events, such as hurricanes. Growth is likely to pick up in a subsequent year due both to rebuilding efforts and a depressed starting point. This does not imply such events are not damaging. However, it does mean that, almost mechanically, output losses are unlikely to translate into a sustained negative growth effect.

sectors (for our primary three bin model with $E = 0, E \in [1, 8]$, and $E \geq 9$). Those results are consistent with the much larger effect sizes we find in agriculture: properly accounting for heat waves (in a strict sense with $\gamma = 366$) is more important in that sector.

Second, we consider an extension of the standard practice of counting days in which the temperature falls into different bins. That practice could be extended to examine the effects of different sequences of exposure to temperatures, for example, counting the number of times during a year that a day between 90 and 100F was preceded by 2 days between 80 and 90F. To allow for full flexibility, for B temperature bins and D -day sequences, we would need to estimate effects of $B^D - 1$ different sequences (after omitting a reference sequence), which quickly becomes intractable. Further interpretation in such models is challenging, as changing a single day's temperature affects many adjacent sequences simultaneously (see Online Appendix A.8).

Still, we use this sequence-based approach in two ways. First, we use the low-dimensional case of $B = 2$ and $D = 6$ to test whether the effects of heat are non-additive across days. Specifically, we count 6-day sequences defined by whether the temperature on each day is above or below our 1.5 SD threshold (e.g. 000001 is a hot day preceded by five cooler days). We then compute the annual aggregates of distributed lags (i.e. we count days for which the d th lag ($d \in [0, 5]$) was abnormally hot). Including those quantities on the right side of equation (7) means the sequence coefficients represent interactive effects of heat, holding the overall incidence of hot days fixed. For all output measures, we are able to reject the joint null hypothesis ($p < 10^{-11}$) that all sequence coefficients are zero, implying that the sequencing of exposure to heat matters.

We can also interpret our binary accumulation model as a special case of this bin-based sequence approach. Our binary accumulation model uses $B = 2$ for time- and location-specific bin definitions and various values of D . It imposes further assumptions, requiring that only the sequence with D days in the abnormally-hot temperature bin (i.e. sequence 111111 for $D = 6$) has a nonzero coefficient. The “partial relief” model described above with finite γ imposes a different restriction on the sequence coefficients, effectively grouping different sequences into a single “treated” set (e.g. sequences 11101 and 11011 are assumed to have the same effect). Other groupings are of course possible, and simply impose other coefficient restrictions to ask different questions. See Online Appendix A.8 for an example.

7. Discussion and Conclusion

Heat waves have attracted substantial attention in the popular press, often in the form of blame for mortality and health effects, damage to crops, and strain on energy systems. However, heat waves have received surprisingly little direct attention in the economics literature, which instead tends to focus on the effects of days or long periods of time that are hot on average. With this paper, we aim to fill that gap, focusing on both the measurement of heat waves and their impact on economic output at the country level.

We find that heat waves depress per-capita economic output in the agricultural sector, with long or severe waves also impacting non-agricultural and overall output. In light of past and projected increases in heat wave frequency and severity, these findings suggest that damages from climate change may be larger than are suggested by standard empirical approaches and integrated assessment models. In turn, current estimates of the SCC, which focus on damages from shifts in mean temperatures, are quite likely too low. In sum, we take these results to imply that temperature may impact economic activity in ways that have not been adequately captured until now.

To the extent that climate-induced shifts in heat wave occurrence and intensity spur changes in anticipated risks, adaptation could mitigate the effects of future heat waves. Further, adaptation to expected climate-induced warming could also affect the impact of heat waves through, for example, purchase of air conditioners. Those adaptations are likely to mitigate the effects of waves by reducing exposure to high heat, but could plausibly also have perverse effects as well (e.g. if electrical grids fail under the load from additional air conditioning use during a heat wave).

Adapting to heat waves poses distinct challenges. First, heat waves are rare, and while overall rates of incidence are projected to increase, specific events are often difficult to predict more than 1–2 weeks in advance. Expectations based on such short-term forecasts can certainly inform minor adjustments (e.g. purchasing a fan, going out of town for a few days, or slightly shifting a production schedule). However, costlier changes that must be undertaken farther in advance, such as planting more heat-tolerant crops, upgrading building insulation, or relocating, carry highly uncertain benefits. Given difficulties many individuals have with assessing the probability of and making decisions about rare events (Kahneman and Tversky 1979), adaptation to heat waves is likely to be far from optimal. Evidence on adaptation to other rare weather phenomena (e.g. hurricanes and floods) is consistent with this claim: protective investment is often inefficiently low and strongly tied to recency of exposure (Kunreuther and Slovic 1978; Meyer 2012). In contrast, adaptation to more predictable temperature changes, such as gradual warming trends, is not plagued by these rarity-induced challenges.

While specific heat waves are difficult to anticipate far in advance, Figure 1 suggests that El Niño events may provide useful information about increased overall likelihood (Online Appendix Section A.11).²³ Specifically, in our sample, the incidence of heat waves is significantly higher in the year following an El Niño event, especially for countries in which average temperature is strongly linked to ENSO (“teleconnected” countries; see e.g. Hsiang, Meng, and Cane (2011)). These patterns are consistent with prior work suggesting a link between ENSO and heat wave formation (e.g. Keellings and Waylen 2015; Luo and Lau 2019). Because advance prediction of El Niño events is improving, established ENSO indices could be useful when making heat wave adaptation decisions. Investigating how best to leverage ENSO information in adapting to heat waves—and whether individuals or businesses already successfully do so—appears a promising area for future work.

23. We thank an anonymous reviewer for highlighting this connection.

While we focus on aggregate economic output for our empirical exercise, our approach to measuring heat waves may be relevant in a variety of other empirical settings. In particular, we demonstrate that our heat-wave metrics generally display low correlation with standard bin-based or moment-based measures of temperature, implying they carry new information. That new information may be an important determinant of micro-level outcomes and could potentially be tied to more specific mechanisms (e.g. heat-related stress) through which heat waves impact economic activity. Similarly, while we find no evidence of long-run effects of heat waves at the scale of national economies, micro-level data could reveal different patterns if, for example, heat waves have lasting, irreversible impacts on human health or infrastructure. Because a number of studies suggest heat-wave incidence is likely to rise in the future, the importance of understanding those effects and mechanisms is only likely to increase.

References

- Angrist, Joshua D. and Jörn-Steffen Pischke (2008). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Auffhammer, Maximilian and Erin T. Mansur (2014). "Measuring Climatic Impacts on Energy Consumption: A Review of the Empirical Literature." *Energy Economics*, 46, 522–530.
- Auffhammer, Maximilian and Wolfram Schlenker (2014). "Empirical Studies on Agricultural Impacts and Adaptation." *Energy Economics*, 46, 555–561.
- Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro (2016). "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship Over the Twentieth Century." *Journal of Political Economy*, 124, 105–159.
- Bita, Craita and Tom Gerats (2013). "Plant Tolerance to High Temperature in a Changing Environment: Scientific Fundamentals and Production of Heat Stress-Tolerant Crops." *Frontiers in Plant Science*, 4, 273.
- Burke, Marshall and Kyle Emerick (2016). "Adaptation to Climate Change: Evidence from US Agriculture." *American Economic Journal: Economic Policy*, 8, 106–40.
- Burke, Marshall, Solomon M. Hsiang, and Edward Miguel (2015). "Global Non-Linear Effect of Temperature on Economic Production." *Nature*, 527, 235.
- Carleton, Tamma A. (2017). "Crop-Damaging Temperatures Increase Suicide Rates in India." *Proceedings of the National Academy of Sciences*, 114, 8746–8751.
- Carleton, Tamma A. and Solomon M. Hsiang (2016). "Social and Economic Impacts of Climate." *Science*, 353, aad9837.
- Charng, Yee-yung, Hsiang-chin Liu, Nai-yu Liu, Wen-tzu Chi, Chun-neng Wang, Shih-hsun Chang, and Tsu-tsuen Wang (2007). "A Heat-Inducible Transcription Factor, HsfA2, is Required for Extension of Acquired Thermotolerance in *Arabidopsis*." *Plant Physiology*, 143, 251–262.
- CIESIN (2016). *Gridded Population of the World, Version 4 (GPWv4): Population Density*. Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC), <http://dx.doi.org/10.7927/H4NP22DQ>, accessed 5 February 2018.
- Conley, Timothy G. (1999). "GMM Estimation with Cross Sectional Dependence." *Journal of Econometrics*, 92, 1–45.
- Coumou, Dim, Giorgia Di Capua, Steve Vavrus, Lei Wang, and Simon Wang (2018). "The Influence of Arctic Amplification on Mid-Latitude Summer Circulation." *Nature Communications*, 9, 2959.
- Cowan, Tim, Ariaan Purich, Sarah Perkins, Alexandre Pezza, Ghyslaine Boschat, and Katherine Sadler (2014). "More Frequent, Longer, and Hotter Heat Waves for Australia in the Twenty-First Century." *Journal of Climate*, 27, 5851–5871.
- Das, Saudamini and Stephen C. Smith (2012). "Awareness as an Adaptation Strategy for Reducing Mortality from Heat Waves: Evidence from a Disaster Risk Management Program in India." *Climate Change Economics*, 3, 1250010.

- 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, 66–95.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (2014). "What do We Learn from the Weather? The New Climate-Economy Literature." *Journal of Economic Literature*, 52, 740–98.
- Deryugina, Tatyana and Solomon M. Hsiang (2014). "Does the Environment Still Matter? Daily Temperature and Income in the United States." NBER Working Paper No. 20750.
- Gasparrini, Antonio and Ben Armstrong (2011). "The Impact of Heat Waves on Mortality." *Epidemiology*, 22, 68.
- Graff Zivin, Joshua and Matthew Neidell (2012). "The Impact of Pollution on Worker Productivity." *American Economic Review*, 102(7), 3652–73.
- Graff Zivin, Joshua and Matthew Neidell (2014). "Temperature and the Allocation of Time: Implications for Climate Change." *Journal of Labor Economics*, 32, 1–26.
- Heyes, Anthony, Matthew Neidell, and Soodeh Saberian (2016). "The Effect of Air Pollution on Investor Behavior: Evidence from the S&P 500." NBER Working Paper No. 22753.
- Hsiang, Solomon (2016). "Climate Econometrics." *Annual Review of Resource Economics*, 8, 43–75.
- Hsiang, Solomon M. and Amir S. Jina (2014). "The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence from 6,700 Cyclones." NBER Working Paper No. 20352.
- Hsiang, Solomon M., Kyle C. Meng, and Mark A. Cane (2011). "Civil Conflicts are Associated with the Global Climate." *Nature*, 476, 438.
- Kahneman, Daniel and Amos Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47, 263–292.
- Keelings, David and Peter Waylen (2015). "Investigating Teleconnection Drivers of Bivariate Heat Waves in Florida Using Extreme Value Analysis." *Climate Dynamics*, 44, 3383–3391.
- Kovats, R. Sari and Shakoor Hajat (2008). "Heat Stress and Public Health: A Critical Review." *Annual Review of Public Health*, 29, 41–55.
- Kunreuther, Howard and Paul Slovic (1978). "Economics, Psychology, and Protective Behavior." *American Economic Review*, 68(2), 64–69.
- Luo, Ming and Ngar-Cheung Lau (2019). "Amplifying Effect of ENSO on Heat Waves in China." *Climate Dynamics*, 52, 3277–3289.
- Mann, Michael E., Stefan Rahmstorf, Kai Kornhuber, Byron A. Steinman, Sonya K. Miller, Stefan Petri, and Dim Coumou (2018). "Projected Changes in Persistent Extreme Summer Weather Events: The Role of Quasi-Resonant Amplification." *Science Advances*, 4, eaat3272.
- Massetti, Emanuele and Robert Mendelsohn (2015). "How do Heat Waves, Cold Waves, Droughts, Hail and Tornadoes Affect US Agriculture?" CMCC Research Paper (RP0271).
- Matsuura, Kenji and Cort J. Willmott (2012). *Terrestrial Precipitation: 1900–2017 Gridded Monthly Time Series*. Version 5.01. <http://climate.geog.udel.edu/~climate/html-pages/download.html#P2017>, accessed 3 December 2018.
- Meehl, Gerald A. and Claudia Tebaldi (2004). "More Intense, More Frequent, and Longer Lasting Heat Waves in the 21st Century." *Science*, 305, 994–997.
- Meyer, Robert J. (2012). "Failing to Learn from Experience about Catastrophes: The Case of Hurricane Preparedness." *Journal of Risk and Uncertainty*, 45, 25–50.
- Nairn, John and Robert Fawcett (2014). "The Excess Heat Factor: A Metric for Heatwave Intensity and its Use in Classifying Heatwave Severity." *International Journal of Environmental Research and Public Health*, 12, 227–253.
- National Academies of Sciences, Engineering and Medicine (2017). *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide*. The National Academies Press, Washington, DC, <https://www.nap.edu/catalog/24651/valuing-climate-damages-updating-estimation-of-the-social-cost-of>, accessed 1 December 2020.
- Ramankutty, Navin, Amato T. Evan, Chad Monfreda, and Jonathan A. Foley (2008). "Farming the Planet: 1. Geographic Distribution of Global Agricultural Lands in the Year 2000." *Global Biogeochemical Cycles*, 22, GB1003, doi:10.1029/2007GB002952.

- Rocklöv, Joacim and Bertil Forsberg (2008). "The Effect of Temperature on Mortality in Stockholm 1998–2003: A Study of Lag Structures and Heatwave Effects." *Scandinavian Journal of Public Health*, 36, 516–523.
- Russo, Simone, Alessandro Dosio, Rune G. Graversen, Jana Sillmann, Hugo Carrao, Martha B. Dunbar, Andrew Singleton, Paolo Montagna, Paulo Barbola, and Jürgen V. Vogt (2014). "Magnitude of Extreme Heat Waves in Present Climate and Their Projection in A Warming World." *Journal of Geophysical Research: Atmospheres*, 119, 12–500.
- 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, 113–125.
- Schlenker, Wolfram and Michael J. Roberts (2009). "Nonlinear Temperature Effects Indicate Severe Damages to US Crop Yields Under Climate Change." *Proceedings of the National Academy of Sciences*, 106, 15594–15598.
- Somanathan, E., Rohini Somanathan, Anant Sudarshan Meenu Tewari, et al. (2015). "The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing." Working Paper No. 244, Indian Statistical Institute: New Delhi, India.
- Thrasher, Bridget, Edwin P. Maurer, Philip B. Duffy, and Colin McKellar (2012). "Bias Correcting Climate Model Simulated Daily Temperature Extremes with Quantile Mapping." *Hydrology and Earth Systems Science*, 16, 3309–3314.
- Wahid, Abdul, Saddia Gelani, M. Ashraf, and Majid R. Foolad (2007). "Heat Tolerance in Plants: An Overview." *Environmental and Experimental Botany*, 61, 199–223.

Supplementary Data

Supplementary data are available at [JEEASN](#) online.