

Harvard Environmental Economics Program

DEVELOPING INNOVATIVE ANSWERS TO TODAY'S COMPLEX ENVIRONMENTAL CHALLENGES

May 2017
Discussion Paper 17-73

Will We Adapt? Labor Productivity and Adaptation to Climate Change*

Jisung Park

Ph.D. 2017, Harvard University, Economics

*This paper won the Harvard Environmental Economics Program's 2017
Ana Aguado Prize for the Best Paper by a Doctoral Student.



HARVARD Kennedy School
JOHN F. KENNEDY SCHOOL OF GOVERNMENT

heep@harvard.edu
<http://heep.hks.harvard.edu>

Will We Adapt? Labor Productivity and Adaptation to Climate Change*

Jisung Park

Harvard University

**This paper won the Harvard Environmental Economics Program's 2017 Ana Aguado Prize for the Best Paper by a Doctoral Student. Before passing away in October 2016, Ana Aguado, for whom the doctoral prize is named, was a leader in European energy policy and business communities for twenty years. She was Secretary General of the European Distribution System Operators' Association for Smart Grids from 2014 to 2016. Before holding that post, Ms. Aguado was Chief Executive Officer of Friends of the Supergrid, dedicated to building an integrated power grid in Europe.*

The Harvard Environmental Economics Program

The Harvard Environmental Economics Program (HEEP) develops innovative answers to today's complex environmental issues, by providing a venue to bring together faculty and graduate students from across Harvard University engaged in research, teaching, and outreach in environmental, energy, and natural resource economics and related public policy. The program sponsors research projects, convenes workshops, and supports graduate education to further understanding of critical issues in environmental, natural resource, and energy economics and policy around the world. For more information, see the Program's website: <http://heep.hks.harvard.edu>.

Acknowledgements

The Enel Endowment for Environmental Economics at Harvard University provides major support for HEEP. The Endowment was established in February 2007 by a generous capital gift from Enel SpA, a progressive Italian corporation involved in energy production worldwide. HEEP receives additional support from the affiliated Enel Foundation.

HEEP also receives support from BP, Chevron Services Company, and Shell. HEEP enjoys an institutional home in and support from the Mossavar-Rahmani Center for Business and Government at the Harvard Kennedy School. HEEP collaborates closely with the Harvard University Center for the Environment (HUCE). The Center has provided generous material support, and a number of HUCE's Environmental Fellows and Visiting Scholars have made intellectual contributions to HEEP.

HEEP and the closely-affiliated Harvard Project on Climate Agreements are grateful for additional support from the Harvard University Climate Change Solutions Fund; the Harvard Global Institute; the Belfer Center for Science and International Affairs and the Ash Center for Democratic Governance and Innovation—both located at the Harvard Kennedy School; Christopher P. Kaneb (Harvard AB 1990); and the International Emissions Trading Association (IETA).

Previous sponsors include: the Alfred P. Sloan Foundation, the AVINA Foundation, Bank of America, Castleon Commodities International LLC, ClimateWorks Foundation, Doris Duke Charitable Foundation, Duke Energy Corporation, the James M. and Cathleen D. Stone Foundation, the National Science Foundation, the U.S. Department of Energy, and the U.S. Environmental Protection Agency.

Citation Information

Park, Jisung. "Will We Adapt? Labor Productivity and Adaptation to Climate Change." Discussion Paper 2017-73. Cambridge, Mass.: Harvard Environmental Economics Program, May 2017.

The views expressed in the Harvard Environmental Economics Program Discussion Paper Series are those of the author(s) and do not necessarily reflect those of the Harvard Kennedy School or of Harvard University. Discussion Papers have not undergone formal review and approval. Such papers are included in this series to elicit feedback and to encourage debate on important public policy challenges. Copyright belongs to the author(s). Papers may be downloaded for personal use only.

Will We Adapt? Labor Productivity and Adaptation to Climate Change

Jisung Park
Harvard University*

April 3, 2017

Abstract

This study explores the for labor-related production impacts of temperature stress both for its own interest and to understand the scope for adaptation to climate change. Focusing on non-agricultural output, I find that hot temperature exerts a significant causal impact on local labor product, with substantially larger effects in highly exposed industries such as construction, manufacturing, and transportation. Places that experience more extreme heat exposure in expectation (e.g. Houston, Orlando) exhibit lower impacts per hot day than cooler regions (e.g. Boston, San Francisco). A year with 10 additional 90°F days would reduce output per capita in highly exposed sectors by -3.5% in counties in the coldest quintile and -1.3%, roughly a third, in the warmest quintile. County-level air-conditioning penetration explains a large proportion of these differences. While these estimates suggest adaptation to heat stress in the long-run, they also imply realistic limits, at least given current technologies.

JEL Codes: D6, N5, Q59, O13, Q1, Q54

Keywords: Climate change, adaptation, economic growth, productivity

*The author would like to thank Lawrence Katz, Andrei Shleifer, Robert Stavins, Joseph Aldy, Raj Chetty, Edward Glaeser, Robert Mendelsohn, Dan Schrag, Peter Huybers, Geoffrey Heal, James Stock, Greg Mankiw, Nathan Hendren, Patrick Behrer, Richard Sweeney and numerous seminar participants at Harvard, Yale, Columbia, Oxford, and Seoul National University for helpful comments and feedback, and Jason Dong for excellent research assistance. All remaining errors are my own. Funding from the Harvard Environmental Economics Program, the National Science Foundation, and the Harvard Climate Change Solutions Fund is gratefully acknowledged.

1 Introduction

Emerging empirical evidence suggests that temperature stress can affect production-related outcomes including labor supply, labor productivity, and total output in the short run, in both developing and developed economies.¹ As Dell, Jones, and Olken (2014) note in a recent review of this rapidly evolving literature, estimates of the labor productivity impacts of heat stress seem to converge to around 1% to 3% normalized decline per °C above room temperature. However, it is unclear to what extent such estimates can currently be used to inform economic policy. In particular, given the gradual and long-term nature of climate change, many have argued that it is critically important to incorporate the potential for adaptation over time in estimating the social cost of carbon (Kahn, 2016).

Will economic agents adapt to future climate change, reducing the realized economic costs of a hotter world? Or will adaptation to climate change be slow, costly, and constrained by practical limits?

So far, evidence for long-run adaptation to changes in climate are mixed, with some analyses suggesting substantial scope for effective adaptation (e.g. Mendelsohn, Nordhaus, and Shaw, 1994; Butler and Huybers, 2013; Barecca et al, 2016) and others finding the opposite (e.g. Annan and Schlenker, 2015; Burke and Emerick, 2016). Moreover, the existing literature on adaptation has focused primarily on agriculture and human health, despite the fact that, as a proportion of total climate damages, labor productivity impacts may exceed all other impacts combined (Heal and Park, 2016; Burke, Hsiang, and Miguel, 2016).

This paper investigates the potential for adaptation to the labor impacts of heat stress. Using a historical panel of weather and payroll data from the United States (1986-2011), I compare the marginal impact of short-run (annual) heat exposure across regions that have experienced different long-term (decadal) climates, under the assumption that agents will optimally adapt to average local climates given sufficiently large production impacts. Across a range of specifications that control for county and year fixed effects, I find evidence of significant adaptation to extreme heat in the long run; places where agents can expect more hot days (e.g. days with maximum temperatures above 90°F) on average experience reduced production impacts per hot day. However, the fact that hot days reduce output even in some of the richest, most well-adapted regions of the United States suggest that certain industries remain susceptible to non-trivial temperature-related productivity losses – at least given existing technologies.

The empirical analysis proceeds in four steps. First, I estimate the causal impact of extreme heat on local non-agricultural production by leveraging quasi-random variation in the number of hot days per year within a given county over time. Using estimation models that control for time-invariant unobservables and smooth trends in payroll at the county level, I find that the average U.S. county experiences a 0.03% reduction in payroll per capita

¹See Zivin and Neidell (2013), Cachone et al (2013), Sudarshan and Tewari (2014), Deryugina and Hsiang (2015), and Burke, Hsiang, Miguel (2016).

during a year with one additional day with maximum temperatures above 90°F.² The estimation strategy relies on the assumption that year-to-year fluctuations in the number of extreme heat days in any given county are uncorrelated with unobserved determinants of county payroll after controlling for time-invariant differences (county fixed effects), trends in economic activity (county-specific polynomial time trends and controls for population), as well as correlated output shocks at the national level (year fixed effects).

Second, to further isolate production impacts of heat exposure operating through labor inputs, I first verify that the main effect is not driven by agricultural yield, and then compare the impacts across sectors that are likely to be more or less exposed to the elements.³ Industries classified by the National Institute for Occupational Safety and Health (NIOSH) as highly exposed – namely, construction, transportation, utilities, manufacturing, and mining – experience markedly higher impacts than relatively insulated ones such as education or financial services.⁴

Third, I compare the marginal impact of an additional hot (90°F+) day on output in highly exposed sectors across counties with different average climates to estimate the scope for long-run adaptation. The intuition here is as follows. Assuming producers have an incentive to protect labor inputs from heat-related production impacts, we would expect them to invest in long-run adaptive capital (e.g. air conditioning) to the point where the expected payoffs over time equal the total costs of investment. Given optimizing agents, we would expect the realized marginal impact of an additional hot day in a place like Houston, which experiences 93 days above 90°F per year on average, to be different from the marginal impact in a place like Boston, which experiences only 9 such days per year, due to the fact that the expected benefits of air conditioning in Boston are lower given the cooler average climate.⁵ To the extent that Bostonians and Houstonians are optimally adapted to their current (historical) climates, one might interpret the differences in marginal damage coefficients as an approximation of the potential scope for adaptation to global warming in the long run, at least given existing technologies.

I find that very hot places (e.g. Houston, Orlando) are significantly better adapted to the production impacts of heat stress than colder or milder ones (e.g. Boston, Minneapolis, San Francisco). Regions in the 1st and 5th quintiles of the 90°F+ day distribution — which feature, on average, 3 and 83 such days per year — suffer short-run impacts of -0.35 percentage points and -0.13 percentage points per additional day above 90°F respectively.

²I focus on the impact of heat stress, though controls for days with mild and cold temperatures are included non-parametrically, in addition to controls for average as well as extreme precipitation. In most regions, cold days do not seem to have a significant impact on non-agricultural production, consistent with Deryugina and Hsiang (2015).

³I also attempt to isolate impacts on industries that are more or less susceptible to temperature-driven changes in short-run demand; that is, supply-side production impacts on payroll as opposed to effects arising from changes in demand that are correlated with hot temperature.

⁴Payroll in healthcare-related industries increases slightly in years with more hot days, consistent with the documented relationship between extreme heat and human health (Zivin and Schrader, 2016; Barecca et al, 2016).

⁵Note that this may be true in expectation despite the fact that, on any given unusually hot year, Boston-based producers may suffer large negative impacts from heat exposure.

The short run impact of an additional hot day falls monotonically as one moves to hotter regions within the U.S., suggesting that optimizing agents do in fact respond to persistent temperature extremes.

Fourth, I explore the role of air conditioning in mitigating the adverse production impacts of extreme heat, using a newly constructed panel of residential AC penetration at the county-year level.⁶ I find that a significant proportion of the difference in marginal impacts of heat exposure can be explained by the spread of air conditioning, consistent with recent findings in the context of adaptation to heat-related health impacts (Barecca et al, 2016). Because changes in AC penetration are not experimental, I cannot rule out the possibility that the effects documented here are driven by correlated changes in other unobserved variables. However, AC penetration does not seem to affect the production impacts of colder days, suggesting that the adoption of AC is not coincident with factors that determine overall payroll.

These results build on an extensive literature on the economic impacts of climate change, reviewed by Tol (2009), as well as the growing literature on the welfare impacts of climate change arising from direct heat exposure (Dell, Jones, and Olken, 2014; Heal and Park, 2016). The estimates of temperature-driven economic impacts are broadly consistent with prior results, including Hsiang (2010), Dell, Jones, and Olken (2012), and Deryugina and Hsiang (2015). However, they also imply that, contrary to previous suggestions that developed economies are well-insulated from climate damages, even the world’s wealthiest economies are currently subject to non-trivial weather-related output losses - impacts which may be exacerbated by future climate change.⁷

These findings also contribute to a growing literature on adaptation to environmental change (Mendelsohn, Nordhaus, Shaw, 1994; Hornbeck, 2012; Burke and Emerick, 2016). The method explored here — of leveraging the spatial gradient in temperature sensitivity and the degree of climate adaptation that this implies — builds upon work by Butler and Huybers (2013) in agriculture and Barecca et al (2016) in health, extending it to labor-related settings. One objective of these studies has been to allow researchers to eventually link the econometrically well-identified studies of weather-driven output shocks to the historically more simulation-based estimates of the social costs of carbon (e.g. Stern, 2008; Hope, 2009; Nordhaus, 2010), an objective that this paper shares.

The rest of the paper is organized as follows. Section 2 provides some background information and a simple model that guides the empirical analysis. Section 3 presents the data and summary statistics. Section 4 presents the empirical strategy, and Section 5 presents the main empirical findings. Section 6 discusses and concludes.

⁶Given the focus on production impacts, commercial and/or industrial AC data would be the ideal measure. Such data was not available. The data that is available shows a tight correlation between residential and commercial AC penetration across regions, suggesting that residential AC penetration may provide a good proxy for average AC penetration in production-related sectors for a given region.

⁷Some studies have suggested that developed economies may even benefit from moderate amounts of warming (Tol, 2009; Mendelsohn, 1994).

2 Background and Conceptual Framework

2.1 The Welfare Impacts of Heat Exposure

Recent empirical work finds strong evidence for a causal impact of short-run heat exposure on economically relevant outcomes, including human health, labor productivity, and labor supply. For instance, Deschenes and Greenstone (2011) find that an additional day above 90°F leads to a 0.11% increase in annual mortality in the United States, controlling for location-specific characteristics and the potential for harvesting. In the context of labor productivity, Cachon et al (2012) document significant negative impacts of extreme heat on automobile plant output, controlling for plant-specific productivity and seasonality in production. They find that a week with six or more days above 90°F reduces output that week by 8% on average.⁸ Graff Zivin and Neidell (2013) document substantial contractions in labor supply on hot days in those US industries with high exposure to extreme temperature and weather shocks. They find that, for highly exposed occupations (e.g. construction), days with temperature above 100°F (37°C) lead to 23% lower labor supply than temperatures between 77°-80°F (25°-27°C). These studies, and the longstanding laboratory literature on temperature and task productivity they build upon, form the basis for exploring adaptation to heat stress in the context of non-agricultural production activities.

2.2 Evidence for Adaptation to Direct Heat Stress

How quickly and effectively economic agents can adjust to changes in their environment is a question of central relevance for economic theory as well as economic policy (Samuelson, 1947; Viner, 1958; Mendelsohn, 1994; Davis and Weinstein, 2002; Cutler, Miller, Norton, 2007; Hornbeck, 2012; Burke and Emerick, 2016). At the most general level, economists have debated this issue theoretically since at least Samuelson (1947), who suggested the LeChatelier principle: that longer time horizons will allow for greater margins of adjustment to any given shock or change in the economic environment.

Estimating adaptive responses to environmental changes in the long run is especially important in the context of climate change, which will take place over the span of multiple decades. Despite a rapidly evolving literature that documents a statistically robust relationship between short-run weather variation (e.g. temperature and rainfall shocks) and economic variables of interest (e.g. mortality, labor productivity, conflict, exports), it remains unclear whether these *weather-sensitivities* are reflective of long-run *climate sensitivity* of social welfare, mainly due to the possibility of adaptation.

Many recent studies that utilize short-run weather fluctuations to identify causal impacts of temperature shocks then combine these coefficient estimates with climate model projections to estimate the expected costs of future climate change. What might be some of the limitations this approach? Suppose the temperature-output response functions with

⁸Sudarshan et al (2014) find similar plant-level productivity declines among Indian manufacturers, even when controlling for region, firm, and individual-specific factors.

and without long-run adaptation are as depicted in Figure 1. If what one is interested in from a policy standpoint is the true long-run social costs of climate change, $(V_0 - V_1)$, then estimating this using short run panel impacts, $(V_0 - V_2)$, might overstate damages by $(V_1 - V_2)$, which is the extent of adaptation which occurs over the long run.

In particular, the realized welfare costs of climate change will be highly sensitive to the scope, speed, and adjustment costs associated with adaptation in the long-run. This paper focuses on the first of these three important parameters.

How important is adaptation for climate policy? Generally speaking, one can imagine three stylized possibilities. First, adaptive adjustments may be effective at reducing climate impacts, and occur quickly and at low cost, in which case using short-run weather sensitivities to estimate long-run climate damages would overstate the urgency of public policy in addressing climate change. Alternatively, it may be the case that adaptive investments occur slowly, are prohibitively costly, and/or suffer from various market failures or collective action problems. This would suggest that economic damages under climate change would likely be large, implying a more substantial role for public policy in addressing future climate threats. A third possibility is that, regardless of the potential effectiveness of some adaptive investments, the set of adaptation options actually shrinks in the long run, due to the depletion of finite resource stocks such as fossil aquifers or ecological buffer capacity.

The economic literature on adaptation has to date focused primarily on agricultural yield (Mendelsohn et al, 1994; Mendelsohn et al, 2000; Schlenker and Roberts, 2011; Butler and Huybers, 2012; Burke and Emerick, 2016). The evidence is mixed, with some studies suggesting substantial scope for adaptation to heat stress, and others finding weak or inconclusive evidence that farmers adapt to changes climate.⁹

Recently, Barecca et al (2016) find evidence for adaptation in the context of health responses to temperature shocks. They find that the mortality impacts of heat stress in the United States, which are most acute in months with days above 90°F, declined rapidly over the twentieth century: by roughly 75 percent, most of it occurring after 1960.¹⁰ Using state-level air conditioning penetration data, they find that the vast majority of this decline can be explained by adoption of air conditioning as opposed to electrification or the number of physicians in the state.

This paper addresses the prospect of adaptation to the impacts of heat stress on labor inputs. The intention is to include all possible economic sectors that are subject to temperature-related production impacts arising from thermal stress of the human body — including labor supply, task productivity, and direct disutility — and to estimate the extent of potential future adaptation using observed (as opposed to simulated) data. In

⁹A limited number of global and regional adaptation cost assessments exhibit a large range and have been completed mostly for developing countries, with the most recent and most comprehensive to date global adaptation costs range from \$70 to more than \$100 billion annually by 2050 (World Bank, 2010). But the quantity and quality of local studies varies by sector with more treatment of adaptation in coastal zones and agriculture (Agrawala and Fankhauser, 2008).

¹⁰Barecca et al (2016) use daily mean temperatures, and find that days with mean temperatures above 80°F cause the majority of temperature-related mortality.

most OECD countries, non-agricultural output accounts for over 95% of total income, which arguably makes adaptation in the context of non-agricultural production of central importance in estimating the true social costs of carbon.

2.3 A Simple Model of Adaptation to Labor Impacts of Temperature Stress

This section outlines a simple model of adaptation to the production impacts of temperature stress, beginning with a brief discussion of key temporal dimensions of adaptation decisions.

The types of adaptation mechanisms available will depend crucially on the time-frame of interest. For instance, in the very short run, individuals may adapt by adjusting labor supply, either on the intensive margin at the daily level, with individuals choosing to work more or less hours or shifting the timing of work hours during a given day, or on extensive margin at the daily level, choosing not to work at all if conditions are bad enough (Zivin and Neidell, 2014).¹¹ Individuals may also adjust exertion levels (labor effort), or engage in other forms of defensive behavior (e.g. wearing different clothing) without changing labor supply (Park and Heal, 2013). In the longer run, persistent temperature shocks may lead workers to change occupations, migrate to a more hospitable climate, or exit the labor force completely due to health concerns or disamenity costs. Presumably, these latter adjustments entail substantially higher pecuniary and non-pecuniary costs, and would only be justified under more extreme levels of environmental stress.

Similarly, flow expenditures on heating and cooling may in most cases be easily adjusted in the short run, but changes in the stock of heating and cooling equipment – for instance, upgrading an air conditioner from window unit to central AC, or retrofitting a home with better insulation – may require longer time horizons.

It is also possible to draw a distinction between secular (private) and directed (public) adaptation responses: that is, to differentiate between those adaptation mechanisms that one would expect to occur naturally in a market economy as a result of changing climates or incomes, and those that would not occur due to important market failures (Table 1 provides a heuristic of adaptation mechanisms by type). For the purposes of this study, I will refer to the former class of adaptation mechanisms as *secular adaptation*, the latter as *directed adaptation*, following Agrawala and Fankhauser (2008). I discuss potential market failures in adaptation investment in the Appendix.

The empirical strategy employed in this paper takes a revealed preference approach to inferring the extent of adaptation to climate stress, and thus makes the conservative assumption that what is measured econometrically encompasses some combination of secular and directed adaptation responses that represent the optimal adaptation bundle. Specifically, comparing the realized impacts of temperature stress on output/productivity net of short run adaptations *within each region* with the impacts of temperature stress given

¹¹The environmental and health literatures typically refer to this as “avoidance behavior”.

different levels of long run adaptation *across regions* allows the econometrician to estimate the expected extent of adaptation in the long run, subject to some simplifying assumptions regarding the availability of adaptation strategies (uniform across regions) and lack of barriers to spatial equilibrium in observed labor markets.¹²

2.4 Production Impacts of Heat Stress

To motivate the empirical strategy, I outline a partial equilibrium model of local adaptive investment in response to the production impacts of heat stress.

Define production-relevant temperature stress, T^E , as a measure of extreme heat. For instance, this could be the number of extreme heat days per year, analogous to the concept of killing degree days in the agricultural literature. T^E is a random variable, the historical distribution of which is a reflection of average climate in that area. The existing literature suggests that T^E can reduce human task productivity (i.e. labor productivity) and may affect the direct utility of workers. Let us make the simplifying assumption that extreme heat does not significantly affect the productivity of non-labor inputs (e.g. the productivity of capital).¹³

Consider the production function $Y(A, L)$, which take as inputs labor productivity (A), and labor supply (L), where labor supply includes both dimensions of hours and effort. While the focus here is on labor, it is worth noting that a possible adaptive response to heat stress may be to adjust capital-labor ratios of production, depending on which factor is more temperature sensitive. As Kahn (2016) and Heal and Park (2016) point out, estimating such responses constitutes an important area for future research.

Allowing labor supply and productivity to depend on temperature means that output is a function of experienced temperature:

$$Y(A, L) = Y(A(T^E), L(T^E))$$

We abstract away from principal-agent problems or labor market frictions, such that the revenue impact of a productivity shock is completely internalized.¹⁴ Workers maximize utility, $U(Y, L, T^E)$, which is an increasing function of production (income), and decreasing in labor supply, labor effort, and temperature stress, which causes direct disutility ($\frac{\partial U}{\partial T^E} < 0$).

The task productivity literature suggests that physical and cognitive task productivity falls with extreme temperature — both heat and cold. Here, we focus on the hot end of

¹²Put a different way, by leveraging daily weather data combined with cross-sectional variation in panel estimates, I estimate the gap between climate sensitivity inclusive of 1) short run (intra-annual) adaptation mechanisms, and 2) long run (decadal) adaptation mechanisms.

¹³It is possible that the effectiveness of physical capital may be sensitive to extreme heat. For instance, heat rates at power plant are affected by ambient temperature, and electronics are known to malfunction at high temperatures. Whether extreme heat has a first-order effect on capital product is a question that remains yet unresolved.

¹⁴A stylized representation of a production context that has these features may be a family-owned, family-operated business, or a one-man rickshaw operation.

the temperature-task productivity relationship, such that $\frac{\partial A}{\partial T^E} < 0$.¹⁵ Existing studies also suggest that labor supply, defined here as a combination of labor hours and labor effort, reacts negatively to extreme temperature, in part due to direct disutility, in part due to lower productivity: $\frac{\partial L}{\partial T^E} < 0$.¹⁶ It is possible to show that, absent strong income effects, temperature deviations from the thermoregulatory optimum will affect labor hours and labor effort in the same direction (Park and Heal, 2013), such that heat shocks will reduce effective labor product, net of optimizing responses of workers who may reallocate labor effort and hours accordingly.

As such, I assume that $\frac{\partial A}{\partial T^E} < 0$ and $\frac{\partial L}{\partial T^E} < 0$. This implies that extreme heat will reduce total output due to this reduction in total labor product:

$$\frac{dY[A(T^E), L(T^E)]}{dT^E} < 0$$

Importantly, given utility-maximizing workers who have some flexibility in their choice of work hours or effort, realized output fluctuations in response to temperature shocks should be net of adjustments on the labor supply and labor effort margins.

2.4.1 Long-Run Adaptive investments

Suppose firms undertake structural adaptive investments, α , which reduce the negative impact of extreme heat stress by reducing the temperature sensitivity of workers' task productivity:

$$\frac{d^2 A}{dT^E d\alpha} > 0,$$

and/or reducing the temperature sensitivity of labor supply (effort and hours):

$$\frac{d^2 L}{dT^E d\alpha} > 0.$$

In principle, firms might be able to engage in such adaptive investments in both the short and long run: for instance, operating existing window AC units more intensively in response to a few hot days (short run); deciding to install central AC in response to a perceived shift in the long-run climate distribution (long run). Here, I focus on the decision to invest in long-run adaptive capital, which may take the form of structural investments such as centralized cooling systems or cultural capital in the form of procedural norms: for instance, adjusting daily worker schedules to minimize heat stress, as has been documented in many tropical countries. One would expect such long-run investments to be a function of the average climate, $\alpha_{it}(E(T^E))$.

¹⁵The empirical analysis presented below suggests that the labor productivity impacts of extreme cold are relatively small, at least in developed economies such as the United States.

¹⁶Graff Zivin and Neidell (2014) document significant changes in hours worked in response to extreme heat and cold, especially in highly exposed sectors.

Firms will choose to invest in adaptive capital such that the expected marginal benefit associated with additional adaptive investment (in terms of heat-related damages avoided) is equal to the marginal cost. In an unchanging climate, a reasonable proxy for expected benefits would be provided by the average historical incidence of extreme heat events over the period in which the local climate has been observed:

$$E(T^E) \approx \sum_{\tau=1}^t T_{i\tau}^E$$

and the output reductions they have caused, where τ represents the first relevant period. Thus, let us assume that these adaptive investments are sufficiently lumpy so as not to be adjustable in response to acute heat stress (the short run), but rather have been chosen prior to the realization of current extreme heat stress, $T_{i\tau}^E$.

The production function can be written as:

$$Y_{it}(A, L) = Y_{it}(A(T_{it}^E, \alpha_{it}(\sum_{\tau=1}^{t-1} T_{i\tau}^E)), L(T_{it}^E, \alpha_{it}(\sum_{\tau=1}^{t-1} T_{i\tau}^E))),$$

where

$$\begin{aligned} \frac{dA}{dT^E} < 0, \frac{\delta^2 A}{\delta T^E \delta \alpha} > 0; \frac{dA}{dT}(T^E, \alpha) \\ \frac{dL}{dT^E} < 0, \frac{\delta^2 L}{\delta T^E \delta \alpha} > 0; \frac{dL}{dT}(T^E, \alpha) \end{aligned}$$

Output is a function of labor productivity, labor supply, and adaptive capital. Labor productivity and supply at any given point in time will depend not only on the contemporaneous temperature, T_{it} , but also the history of temperature shocks in that location — $\sum_{\tau=1}^{t-1} T_{i\tau}^E$, that is, the local climate — due to the fact that adaptive capital stock will have been chosen to maximize profits subject to the conditions mentioned above.

2.4.2 Application to empirical strategy

The overall effect of adaptive investments will be to reduce the short-run temperature-sensitivity of total output:

$$\frac{d \left| \frac{dY_{it}}{dT_{it}^E} \right|}{d\alpha} < 0.$$

Thus, in the long run, one would expect firms in hotter climates (i=H) to exhibit higher levels of adaptive investment than cooler ones (i=C):

$$\alpha_H > \alpha_C;$$

given

$$E(T_H^E) > E(T_C^E).$$

This paper aims to estimate the production impacts of extreme heat stress, $\frac{dY_{it}}{dT_{it}^E}$, in addition to the expected extent of long run adaptation, $\alpha_H - \alpha_C$, by using differences in realized production impacts across various climate regions:

$$\frac{dY_C}{dT_C^E} - \frac{dY_H}{dT_H^E}.$$

3 Data and Summary Statistics

3.1 County-Level Payroll Data

I use payroll data from the County Business Patterns database from 1986-2012, which records annual and 1st quarter payroll for roughly 3,000 US counties by five-digit NAICS classification. Payroll includes all forms of compensation, including salaries, wages, commissions, dismissal pay, bonuses, vacation allowances, sick-leave pay, and employee contributions to qualified pension plans paid during the year to all employees.¹⁷ County specific payroll data is measured at the annual level.

The choice of payroll as the dependent variable of interest, rather than, for instance, total profits or total income, is motivated by two factors. First, payroll data from the CBP allows one to isolate production impacts on non-agricultural sectors, as well as to distinguish, as I do below, between sectors that are more or less exposed to temperature stress. Second, changes in per capita payroll provide close proxies to changes in total and marginal labor product, separately from changes in capital expenditures. Importantly, payroll is less likely to include direct expenditure on heating or cooling, which may be the case for total income. This means that one is in principle able to estimate the implied marginal benefits of adaptation (in terms of reduced production impacts) separately from the marginal costs.¹⁸

Thus, changes in payroll might be thought of as net fluctuations in the wage bill after firms and individuals each optimize internally, be that in the form of adjustments to labor supply, labor effort, involuntary changes in labor productivity, or short- and long-run investments in adaptive behavior.

¹⁷For corporations, payroll includes amounts paid to officers and executives; for unincorporated businesses, it does not include profit or other compensation of proprietors or partners. Payroll is reported before deductions for social security, income tax, insurance, union dues, etc. This definition of payroll is the same as that used by the Internal Revenue Service (IRS) on Form 941 as taxable Medicare Wages and Tips (even if not subject to income or FICA tax).

¹⁸For instance, if a manufacturing firm pays workers' wages as a function of hours worked and items produced, fluctuations in payroll arising from temperature shocks would reflect changes in labor supply and labor productivity (as well as, in principle, demand for the product itself) which arise in response to heat stress. If, in addition, firms respond by running air conditioning equipment at a higher utilization rate, this added cost would be reflected in lower profits or net income, thus conflating some portion of realized output shocks with short-run flow expenditures on adaptive capital. Payroll, which more closely approximates marginal product of labor than capital, seems less likely to do so.

3.2 Daily Weather Data

County-level payroll data is matched with daily weather data from the PRISM model, which provides temperature and precipitation readings for a 2km x 2km grid of the contiguous United States. Daily max, min, and average temperatures, in addition to precipitation are area-weighted to the county level, and variables containing the number of days with daily maximum temperatures in a series of 10°F bins are constructed by county and year.

Past literature has documented a persistent, non-linear relationship between temperature and economic outcomes, particularly in the context of extreme heat stress (Schlenker and Roberts, 2009; Hsiang, 2010; Deschenes and Greenstone, 2011; Burke and Emerick, 2016; Barecca et al, 2016). Where data has been available, this relationship has been captured using the concept of temperature days: for instance, growing degree days, GDD, in the case of agriculture, which measure the amount of time a crop is exposed to temperatures between a given lower and upper bound, with daily exposures summed over the growing season to ascertain annual growing degree days.

Here, I use the concept of extreme heat days, which are defined as days with daily max temperatures above 90°F, following Deschenes and Greenstone (2011). Deschenes and Greenstone (2011), Hsiang (2010), Sudarshan (2014), find days above 80°F, 85°F or 90°F respectively to be significant heat thresholds that lead to discernible impacts on human performance in field settings.¹⁹ This concept is also analogous also to Killing Degree Days in the agricultural literature, which has a kink point of roughly 77°F, 25°C (Schlenker and Roberts, 2007; Burke and Emerick, 2016).²⁰

3.3 Air Conditioning Data

Air conditioning penetration by county and year is constructed using county level residential AC information from the 1980 decennial census, and combining it with data on changes in residential AC penetration over time by census region from the Energy Information Agencies Residential Energy Consumption (REC) surveys. I use the reported penetration rates in 1980 as a basis and then extrapolate based on the region-level growth rate of central, window and total AC penetration recorded by RECS, which provide penetration rates by region from 1980 to 2009 with a two or three-year frequency. I linearly interpolate growth rates for the missing years and assign counties their corresponding regional growth rate. Using this growth rate and the observed penetration rate in 1980 I create a measure of penetration in every county in each year from 1980 to 2014. I top-code penetration at 100%. Our primary specification uses the penetration rate of total AC but

¹⁹The kink point is lower in lab studies (e.g. Sepannen, 2008). This could be due to the fact that most lab experiments impose something akin to a no-adaptation constraint. Participants are required to concentrate on challenging tasks under temperature stress, without the ability to rest between sessions, adjust physical surroundings, or adapt production techniques.

²⁰I also use alternative measures of temperature shocks, including cooling degree days and average annual temperatures (daytime high temperatures inclusive of humidity), and present these results in the appendix. For the most part, the results are consistent across different measures of temperature, though they are sharpest using the extreme heat day definition.

I conduct the same exercise for central and window AC and estimate all models with all three measures of AC penetration. The results across all three measures are qualitatively similar.

Given the focus on production impacts operating through labor inputs, the ideal measure of AC would include commercial and/or industrial AC penetration. Such data was not available. However, available evidence suggests that commercial and residential AC penetration rates are highly correlated within regions. For instance, according to the EIA, 74% of commercial buildings in the Northeast region had some form of AC as of 2009, while approximately 80% of residential buildings did. In the East South Central region, the commercial and residential penetration rates were 90% and 95% respectively. Based on this observation, and the assumption that local determinants of residential and commercial AC are likely to have shared components (see Biddle, 2008), I take residential AC as a proxy for total AC at the county level.

3.4 Summary Statistics

Over the period 1986 to 2012, the average county in the contiguous United States had an average annual temperature of 54.6°F, and experienced approximately 25 days with temperatures above 90°F per year. This masks tremendous variation across regions. Parts of the Northeast and coastal regions of the West seldom experience any days above 90°F. Seattle and San Francisco experienced fewer than one such day per year on average over the period. In contrast, parts of the South and Southwest regularly get more than 80 such days per year. Figure 2 depicts the average incidence of 80°F+ and 90°F+ days across the country, illustrating this variation graphically.

It is important to note that realized temperatures can vary considerably even within small geographic locales (e.g. counties) depending on elevation, distance to bodies of water, vegetation, or surface albedo. For instance, even within Los Angeles County, the temperature on a given summer day may be 30°F lower in Santa Monica, which is on the coast, than it is in Pasadena, which is farther inland. To the extent that our measures of local temperature are measured with (classical) error, we would expect the estimates of the impact of heat exposure on production to be downward attenuated.

Figures 3, 4, and 5 depict imputed average AC penetration rates across the country in the years 1990, 2000, and 2010 respectively, excluding Alaska and Hawaii. As of 1986, the average residential AC penetration rate across all counties was 58%. By 2010, it had risen to 69%. Once again, there is considerable variation across regions, both in initial levels of AC penetration and rates of uptake over time. For instance, AC penetration in New York City rose rapidly during the period 1986 to 2011, increasing from 55% to 89%. In Grady, GA, AC uptake was much slower, from 62% to 69% over the same period (Figure 6). This is despite the fact that most of Georgia experiences far more heat exposure than New York on average: 70 days above 90°F in Grady as opposed to 14 in New York City. Houston, TX, on the other hand, had close to universal AC even as of 1986.

Running simple OLS regressions in the cross-section suggests a strong correlation between productivity and average climate. Pooling all years in the sample, a region with one more heat day (90°F and above) per year on average features 0.478% lower non-agricultural payroll per capita, controlling for precipitation and snow.²¹ This is consistent with the cross-sectional gradient documented by Acemoglu and Dell (2010), who find a within-country slope of roughly -1% per degree F increase in average annual temperature across municipalities in North and South America.

4 Empirical Framework

I begin by describing the regression models used to estimate the relationship between temperature and local output, as measured by payroll per capita. In each of the analyses presented below, causal identification relies on a framework that leverages quasi-random variation in annual temperature, netting out location-specific factors that may affect production.²²

4.1 Regression Framework

4.1.1 Main Effect

All of the analyses presented below are based on estimating variants of the following equation:

$$\ln(y_{ist}) = \sum_k \beta_k TMAX_{itk} + \pi_1 PREC_{it} + \pi_2 PREC2SD_{it} + \gamma_i + \eta_t + f_{s,i}(YEAR_t) + \epsilon_{ist} \quad (1)$$

where y_{ist} is annual payroll per capita in county i , state s , and year t . $PREC_{it}$ and $PREC2SD_{it}$ represent average annual precipitation and a variable indicating the number of extreme precipitation events in each county year. Extreme events are defined as daily precipitation totals two standard deviations above the county-specific average. The variables γ_i and η_t denote county- and year-fixed effects respectively. γ_i controls for time-invariant unobserved factors that may determine the relative productivity of county i (e.g. human capital). η_t accounts for correlated shocks that are common across the United States (e.g. recession years). $f_{s,i}(YEAR_t)$ represents flexible time trends that are allowed to vary at the state- or county-level, and control for smooth changes in payroll over time as well as the potential for correlation between secular regional productivity

²¹Using annual average temperatures to check consistency with the existing literature, I find that counties with one degree F hotter annual temperatures are associated with -0.924% over the pooled sample. The same coefficients are -2.34%, -1.04%, -0.24% in years 1990, 2000, and 2010 respectively. Moreover, fitting a quadratic specification yields a single-peaked relationship between temperature and implied output, suggesting an optimal temperature zone around 52°F average annual temperature. Note that this is somewhat lower than the physiological optimum implied by the medical and task productivity literature (65°F). Some of this may be due to the fact that average annual temperatures include nighttime low temperatures.

²²For a detailed discussion of how the methods employed here relate to the existing literature on climate adaptation, see Appendix.

trends not accounted for by annual population and year fixed effects.

The variables $TMAX_{itk}$ represent our measures of temperature, which are constructed to capture exposure to the full distribution of temperatures in a given year. The $TMAX_{itk}$ variables are defined as the number of days in a count-year in which the daily maximum temperature is in the k^{th} of 9 temperature bins ranging from 0°-10°F to 90°F and above. In practice, the 70°F-79°F bin is the excluded group, so the coefficients on the other bins are interpreted as the effect of exchanging a day in the 70°F-79°F range with a day in other bins. As noted by Deryugina and Hsiang (2015) and Barecca et al (2016), the primary functional form restriction imposed by this model is that the impact of the daily max temperature on annual payroll is constant within 10°F bin intervals.

I use the number of days above 90°F as the primary indicator of extreme heat. This is motivated by previous studies, which find strong impacts of heat stress on human behavior and task productivity beginning around 85°F, as well as the observation that most productive activity occurs during the daytime, motivating a choice of daily max as opposed to min or mean temperature as the primary measure. Not specifying additional bins above or below this threshold represents an effort to remain as non-parametric as possible while also obtaining estimates that are precise enough to permit meaningful interpretation.²³

In all versions of equation 1, the β_k parameters are identified from inter-annual variation in temperature realizations. It seems difficult to come up with other potential confounders that are not captured by the rich controls above, suggesting the identifying assumptions are likely to be satisfied.

4.1.2 Estimating Labor Impacts

To isolate the impact of temperature on non-agricultural sectors, I subtract agricultural payroll from total annual payroll for each county-year, and run a version of equation 1 that uses log non-agricultural payroll as the dependent variable.

To further isolate the impact on labor inputs, I examine impacts by industry, where j denotes industry classification:

$$\ln(y_{istj}) = \sum_k \beta_k TMAX_{itk} + \pi_1 PREC_{it} + \pi_2 PREC2SD_{it} + \gamma_i + \eta_t + f_{s,i}(YEAR_t) + \epsilon_{istj} \quad (2)$$

Determining ex ante which industries are more or less susceptible to temperature stress in an empirically executable way is not an exact science, in part because CBP payroll data is categorized by two-digit NAICS parent codes as opposed to specific occupations.

²³I focus on the 90°F threshold as a conservative definition of extreme heat, though the results are robust to alternative specifications, including multiple critical temperature-bins and a fully non-parametric temperature bin specification. The main results are also robust to alternative measures of temperature, including those that have been used in the literature in the absence of daily weather data, including annual average temperature and cooling degree days.

Each parent category (e.g. Construction, Retail, Transportation) includes many specific occupations that may feature vastly different working environments. For instance “Transportation” includes “Rail-Track Laying and Maintenance Equipment Operators”, who are likely to work outdoors, as well as “Air Traffic Controllers” who are far less likely to work outdoors.

As a conservative categorization scheme, I follow the National Institute for Occupational Safety and Health’s (NIOSH) classification of “highly exposed” industries: namely, construction, manufacturing, utilities, transportation and mining (I exclude agriculture from all remaining analyses).²⁴ I classify the rest – retail, wholesale, health, education, and finance-insurance-real estate – as “not exposed”. To the extent that the comparison of interest is between highly exposed and non-exposed occupations and this classification only crudely approximates the true subset of exposed occupations, we would expect the analysis to provide an underestimate of the difference, as we would be measuring impacts for air traffic controllers alongside railway repair workers within the same “highly exposed” category, and similarly for occupations that may be more likely to work outdoors in the “non-exposed” category.

As described in greater detail below, demand-side factors may affect our estimation of labor-related production impacts. For instance, hot days may induce greater demand for certain products that are complementary to consumption activities during warm weather (e.g. ice cream). They may also lead to avoidance behavior or adverse health outcomes that directly affect demand for services and thus annual payroll (e.g. emergency room visits to hospitals). I attempt to account for some of these factors by examining specific sectors that are likely to be more or less affected by intra-annual demand-side factors separately as well. If demand-side impacts of heat stress and supply-side production impacts operate in the same direction, our estimates of the production impacts would be biased upward.

4.1.3 Estimating Adaptation

To measure adaptation, I first classify counties according to their average climate. The relevant definition of average climate will depend, in part, on the aspects of the climate distribution that affect the relevant investment decisions. To the extent that the existing literature finds non-linear impacts of extreme heat (as opposed to impacts from shifts in average annual temperatures), one might expect the relevant metric to be the expected number of extreme heat days over time.

In practice, I use various moments of the long-run climate distribution to define “climate”. The preferred specification categorizes counties by the average number of hot days with maximum temperatures above 90°F, though the results are qualitatively similar in

²⁴ “Highly exposed” industries include industries where the work is primarily performed outdoors — agriculture, forestry, fishing, and hunting; construction; mining; and transportation and utilities — as well as manufacturing, where facilities are typically not climate-controlled and the production process often generates considerable heat.

specifications that use a lower temperature threshold (e.g. 80°F) or average annual temperatures. All specifications use averages over the period 1986-2012 for consistency.²⁵

I measure the extent of potential long-run adaptation in two ways. First, I run equation 2 separately by quintile of the historical climate distribution, focusing on impacts for highly exposed industries. Note that many “highly exposed” industries involve work that is both outdoors and indoors, especially in the case of manufacturing. Second, I augment equation 2 by adding interactions of the temperature variables with county-specific measures of long-run climate, for the full sample:

$$\ln(y_{istj}) = \sum_k \beta_k TMAX_{itk} + \sum_k \theta_k^{CL} TMAX_{itk} \times T\bar{M}AX_{i,k=9} + \omega T\bar{M}AX_{i,k=9} + \pi_1 PREC_{it} + \pi_2 PREC2SD_{it} + \gamma_i + \eta_t + f_{s,i}(YEAR_t) + \epsilon_{istj} \quad (3)$$

The coefficients θ_k^{CL} on the interaction term measure whether the effect of an additional day in a given temperature range is affected by the average historical incidence of hot days, relative to the effect of the average historical incidence on a day in the omitted 70°F to 79°F bin. According to the model presented in section II, we would expect places that experience greater heat exposure on average to be better adapted to heat stress, and thus experience lower marginal impacts per hot day. This would result in a positive interaction term for days above 90°F.

4.1.4 Exploring the Role of Air Conditioning

To assess the role of air conditioning in reducing the impact of extreme heat on production, I augment equation 2 with measures of air conditioning penetration. I interact interpolated AC penetration at the county-year level with the temperature variables to estimate the role that AC may have played as a modifier on the effect of hot days on production:²⁶

$$\ln(y_{istj}) = \sum_k \beta_k TMAX_{itk} + \sum_k \theta_k^{AC} TMAX_{itk} \times AC_{it} + \lambda AC_{it} + \pi_1 PREC_{it} + \pi_2 PREC2SD_{it} + \gamma_i + \eta_t + f_{s,i}(YEAR_t) + \epsilon_{istj} \quad (4)$$

The 70°F to 79°F temperature bin is again the excluded group among the k temperature ranges. The interaction term thus measures whether the effect of an additional day in a given temperature range is affected by the average AC penetration rate in that county-year, relative to the effect of the average historical incidence on a day in the omitted 70°F to 79°F bin. The hypothesis is that the coefficients on the interaction terms (θ_k^{AC}) will be positive for hot days (k=9), suggesting that investment in AC mitigates the marginal impact of hot days on production. The interpretation assumption being made in using residential AC is that the determinants of a total AC both across and within counties are

²⁵The primary results are robust to using definitions that use previous weather realizations – for instance, averages over the period 1950 to 1985 – instead.

²⁶I also do the same for average AC for the entire period (1986-2012), as a check against results being driven by the AC interpolation scheme.

similar to the determinants of residential AC over the period 1986-2012.

5 Results

5.1 The Production Impacts of Heat Exposure: County-Level Payroll Per Capita

Figure 7 provides a binned scatterplot that motivates the analyses that follow. It shows the relationship between log payroll per capita and the number of hot days by percentile of the hot day distribution, controlling for average differences across counties and years, as well as the other weather controls and time trends noted in equation 1. It suggests a strong negative relationship between hot days during the year and production that year.

Table 2 presents the results from running versions of equation 1 with state- and county-specific time trends. The dependent variable in this case is non-agricultural payroll per capita. Robust standard errors are clustered at the state by year level to allow for spatial correlation of error terms within a given state and year. The estimates suggest that an additional hot day causes a -0.03% (se=0.007) decline in payroll per capita on average. This means that a year with 10 more hot days results in approximately 0.3% lower payroll per capita for the average U.S. county, or that, in any given year, hot days (of which there are on average 25) reduce total per capita payroll by approximately -0.75% from what would otherwise have been the case were all counties to experience the ideal working temperature year-round.

5.2 Exposed versus Non-Exposed Industries

Figures 8 and 9 present binned scatterplots for highly exposed and non-exposed industries respectively. They suggest more acute impacts in sectors where workers are exposed to the elements. An additional hot day causes a statistically significant -0.11% (se=0.02) decline in payroll per capita in highly exposed industries, as opposed to a statistically insignificant -0.016% (se=0.010) decline in non-exposed industries. In highly exposed sectors, a year with 10 additional hot days reduces labor product by approximately 1.1%. This corresponds to a more than 8-fold difference between exposed and non-exposed sectors. These results are consistent with a story of labor productivity decline due to reductions in cognitive capacity and physical functioning from thermal stress of the human body, as well as shocks arising from reduced concentration and increased mistakes, reduced labor effort, and reduced labor supply.²⁷

²⁷Note that while there may be some bias due to selection in the location of highly exposed industries in the cross section (industries that are highly exposed may choose to locate in locations that are typically colder) our panel approach removes this bias. It is also worth noting that many highly-exposed industries—again think of construction—must take place in places that are both hot and cold. While construction workers in hot places may choose to work differently than those in colder places earlier in the morning for example it is not feasible to not have construction as an industry in Houston. Changes in worker behaviors, meanwhile, are what we would consider structural adaptation, the impact of which we are attempting to

Figures 10, 11, and 12 present analagous binned scatterplots for construction, education, and healthcare sectors respectively. As might be expected, construction payroll declines in year with more hot days, likely due to reduction in labor supply and/or productivity. Payroll in the education sector is unaffected by hot temperature. Healthcare payroll seems to increase slightly during years with more hot days, consistent with the existing literature on mortality impacts of heat stress.

5.3 Evidence for Long-Run Adaptation: Comparing Across Climatic Regions

Results from running the climate-specific regressions by quintile of the average hot day distribution are reported in Table 3. A county in the bottom quintile of the extreme heat day distribution (e.g. San Francisco, Seattle) exhibits a short-run weather sensitivity of approximately -0.35 percentage points ($se = 0.12$) per extreme heat day ($90^{\circ}F+$). A relatively hot county at the top quintile of the US average temperature distribution (e.g. Houston, Orlando) has a short-run weather sensitivity of -0.13 percentage points ($se = 0.05$) per extreme heat day: roughly a third the impact. As columns (2) through (4) suggest, the marginal impact of a hot day seems to decline monotonically as one moves to climates that experience more hot days on average. Column (1) of Table 3 presents the results from running equation 3. The coefficient on the interaction term between the number of hot days in a given year and average climate is positive, suggesting that the impact of a hot day on production declines as one moves to places that are hotter in expectation.

The impact of an additional hot day is roughly 63% smaller in counties in the top quintile of historical extreme heat incidence, compared to counties in the bottom quintile, suggesting substantial scope for adaptation given appropriate investments.²⁸ Whether because of AC or other adaptations, private or public, the same $90^{\circ}F$ day seems to have a very different short run impact in Houston than it might in Boston. While the reduction in temperature sensitivity associated with moving from less to more heat-prone areas is large, it is worth noting that, even in these presumably very well-adapted areas, extreme heat days have statistically significant and economically meaningful impacts on output. These estimates suggest that, at least for highly exposed industries such as manufacturing, construction, or transportation, even those counties in the top quintile of extreme heat exposure suffer routine heat-related output impacts of up to -11% per year, given the high incidence of hot days. This is despite near universal air conditioning in many parts of the US South and Southwest.

The fact that the average realized impacts are not equalized across regions suggests either that adaptation costs are not uniform, and/or that there are non-trivial behavioral barriers or spatial rigidities in production that keep firms from achieving the optimal level of

measure.

²⁸Running the analysis by thirds yields similar results. Both specifications suggest monotonically declining temperature sensitivities as one moves to regions with greater degrees of perennial heat stress.

adaptation implied by the simple, friction-less model presented in section II. For instance, it might be the case that certain outdoor industries like construction or mining feature realistic limits to adaptation (at least given current technologies), but must due to their non-traded nature or geographic restrictions take place in sub-optimal climates. I leave a more careful (structural) analysis of these factors for future research.

5.4 The Role of Air Conditioning

Nearly all households in Houston had AC as of 2009, of which 80% were central AC units. In contrast, only 20% of Massachusetts households had central AC, and 21% did not have air conditioning units altogether.²⁹ Such differences in AC represent but one of a potentially very large number of adaptations that local workers, consumers, and firms have evolved over the years in response to different climates. It is worth noting that, while average AC penetration and incidence of hot days is highly positively correlated, it appears that the relationship is far from uniform, especially in hotter regions (Figure 13). This is consistent with some hot but poor regions such as Grady, GA having experienced slow AC uptake relative to what climatic averages might suggest.

Table 4 presents the results from running augmented versions of equation 2 that interact temperature and AC penetration rates. Columns (2) and (3) present interaction terms between hot temperature and average AC penetration by county (1986-2012). Columns (4) and (5) present interaction terms between hot temperature AC penetration by county-year. The coefficients on the interaction terms between hot temperature and AC penetration are positive, suggesting that having more AC helps protect against the production impacts of hot days. The interaction term between AC and cold days (days with max temperature below freezing) is not significant, further suggesting that the effect is operating through the protecting impact of air conditioning against heat.³⁰

6 Discussion and Conclusion

This paper uses county-level payroll and daily weather data to identify the impact of hot temperature on production, and the potential for adaptation to heat stress in the long run. The findings suggest significant but not unlimited scope for adaptation to climate change in the context of production impacts arising from heat stress of labor inputs, focusing on non-agricultural sectors.

²⁹For average Texas households, 18% of total energy usage is devoted to cooling, compared to 1% for Massachusetts households.

³⁰The use of residential (as opposed to commercial or industrial) AC data may speak to one channel through which exposure to extreme heat has been hypothesized to impact labor productivity. Hot days are normally preceded by hot nights that may make it difficult to sleep. A worker who struggles to sleep the day before a hot day at work might display lower productivity because of lack of sleep, induced by heat, rather than due to contemporaneous exposure to extreme heat at work. If this were the primary mechanism, then one would expect to see an impact from increasing residential AC that makes it easier to sleep on hot nights.

I characterize implied climate adaptation by first estimating the short-run heat-shock sensitivity of local output. For the US as a whole, an additional day with daily max temperatures above 90°F results in a -0.03% reduction in the level of per capita payroll that year. Given well-documented wage rigidities, it seems likely that a non-trivial portion of these impacts are related to, though not entirely explained by, reductions in labor supply and labor productivity.

These effects are non-trivially large. Assuming, perhaps conservatively, that impacts scale linearly with the number of extreme heat days, a year with 10 additional 90°F-or-above days would result in approximately -1.1% lower output per capita in exposed industries for the average US County, and up to 3.5% lower output per capita in milder climates such as the Northeast or Pacific Northwest, which are far less accustomed to such heat exposures. While an unlikely scenario, if the entire country were to experience a year with extreme heat stress corresponding to an average year in Houston, which experiences 92 days per year with daily max temperatures above 90°F, the US economy would experience a -2.7% decline in total output per capita: -10.1% per capita in highly exposed sectors. This study thus lends evidence in support of adding labor productivity impacts into integrated assessment models of climate change, which typically assume total damages on the order of a few percentage points of GDP by 2100.

I find substantial geographic variation in short-run impacts of heat stress, suggesting that adaptation depends in part on the history of weather shocks in a given locality: notably, the average number of hot days (above 90°F). These estimates imply that the productivity impacts in a world where agents engage in no adaptation may be as much as three times as large as one in which all individuals adopt optimal adaptive technologies and norms, though it is unclear how quickly and at what cost such adjustment might occur. Unlike simulation studies which trace the hypothetical costs and benefits of adaptation strategies through particular mechanisms, this analysis empirically estimates the temperature sensitivity of local output and how this sensitivity varies with average local climate. This method has the benefit of not requiring the analyst to simulate all adaptation mechanisms.

While these estimates suggest substantial scope for adaptation in the long run, the fact that temperature shocks exert statistically significant and economically meaningful impacts on labor productivity even in the hottest and presumably well-adapted regions of the United States suggests that there may be realistic limits to adaptation to increased heat stress due to climate change: at least using existing technologies. Given much lower AC penetration rates in much of the developing world (and even in other developed economies in Europe and East Asia), these estimates suggest substantial labor productivity impacts in the medium to long run, even with rapid uptake of AC.

The central methodological message of this paper is that it is possible to approximate the extent of future adaptation by comparing the differences between short run heat-shock sensitivities of local economies that have already optimally adapted to varying levels of

local heat stress, echoing recent work by Barecca et al (2016). The intuition is that the cumulative history of weather shocks in a relatively warm region today may provide a valuable indicator for the extent of long-run adaptive investment that relatively cool regions may eventually undertake in the future, assuming similar availability of adaptive technologies (e.g. air conditioning, alteration of norms around time of work). In other words, cross-sectional gradients in realized output sensitivities should reflect net-of-adaptation values across different climates, an intuition that parallels work by Mendelsohn (1994) and others using the Ricardian method in agricultural contexts, but also addresses critiques regarding causal inference often associated with cross-sectional approaches.

This paper raises important questions for future research. For instance, how rational or forward-looking are agents in making adaptive investments? That is, which climatic mean do they use in making investment decisions? In choosing the HVAC system for a manufacturing plant in Boston, a fully rational investor might make her decision based on some weighted average of existing climate projections published by the IPCC. Given limited bandwidth or lack of information, she may alternatively make a decision based on an intuitive sense of historical climate averages. Whether and to what extent such decisions vary systematically based on education or income may be relevant in assessing the distributional consequences of climate mitigation policy, as well as the potential for welfare-enhancing climate adaptation interventions.

Another set of policy-relevant questions involves the welfare economics of adaptation investment. How much of the relevant adaptive investments will be in the form of private goods, such as home air conditioning, versus local or global public goods, such as workplace norms, electricity infrastructure (e.g. peak grid capacity), or new cooling technologies? Moreover the production impacts documented here imply that it may be possible to uncover adaptation cost functions using observed – as opposed to simulated – data. Though the present analysis does not allow for detailed estimation of the costs associated with such long run adaptations, similar analyses using richer data and/or structural estimation techniques may be able to uncover adaptation cost functions. Given the paucity of reliable adaptation cost estimates despite their policy importance, this seems to be a critical area for future research.

Finally, a natural question that arises is whether the extreme heat impacts and scope for adaptation documented here are reflective of what one might expect in other countries, particularly in the developing world. The substantial heterogeneity in temperature sensitivities within the United States, combined with previous (larger) estimates of labor productivity, mortality, and agricultural output declines due to heat stress in developing countries (Schlenker and Lobell, 2010; Sudarshan et al, 2014) suggests that the long-run impacts of climate change may be more severe for the developing world than previously estimated. It is well-documented that rates of air-conditioning have historically tended to follow income growth quite closely, and have neared saturation in warmer parts of the US (Biddle, 2008; Davis et al, 2015). Based on this relationship and relatively low income lev-

els for many households in warmer parts of South Asia, Latin America, and Sub Saharan Africa, one might infer that these coefficients represent conservative damage estimates for most of the developing world.³¹

More broadly, the fact that the relationship between hot days and output is significant in the United States, one of the world's wealthiest and technologically advanced economies, underscores the climate-dependency of much of economic activity, and suggests furthermore that there may be realistic limits to adaptation driven by rising incomes. Even if developing countries such as India or China were to raise their standard of living to US levels, they may potentially still experience temperature-driven productivity losses of multiple percentage points output per year.

³¹Of course, to the extent that previous findings suggest that poor countries experience markedly different climate impacts compared to rich countries (Dell et al, 2013), one must be careful in extrapolating analyses of rich-countries to assess impacts in poorer regions or at the global level. However, the fact that most global integrated assessment models have historically assumed mildly positive impacts of increased temperatures on rich countries suggests that using data from rich countries to establish a lower bound on labor productivity related climate impacts at the global level would be an important contribution.

References

1. Annan, Francis, and Wolfram Schlenker. "Federal crop insurance and the disincentive to adapt to extreme heat." *The American Economic Review* 105.5 (2015): 262-266.
2. Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro. "Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century." *Journal of Political Economy* 124, no. 1 (2016): 105-159.
3. Biddle, Jeff. "Explaining the spread of residential air conditioning, 1955 to 1980." *Explorations in Economic History* 45.4 (2008): 402-423
4. Burke, Marshall, and Kyle Emerick. "Adaptation to Climate Change: Evidence from US Agriculture." *American Economic Journal: Economic Policy* 8.3 (2016): 106-40.
5. Butler, E. E., & Huybers, P. (2013). "Adaptation of US maize to temperature variations." *Nature Climate Change*, 3(1), 68-72.
6. Cachon, G., Gallino, S., & Olivares, M. (2012). "Severe weather and automobile assembly productivity". The Wharton School, University of Pennsylvania. http://opim.wharton.upenn.edu/~cachon/pdf/weather_1015.pdf.
7. Davis, Donald R., and David E. Weinstein. "Bones, bombs and break points: the geography of economic activity". No. w8517. National Bureau of Economic Research, 2001.
8. Davis, Lucas W., Alan Fuchs, and Paul Gertler. "Cash for coolers: evaluating a large-scale appliance replacement program in Mexico." *American Economic Journal: Economic Policy* 6.4 (2014): 207-238.
9. Dell, Melissa, Benjamin F. Jones and Benjamin A. Olken. "Temperature and Income: Reconciling New Cross-Sectional and Panel Estimates." *American Economic Review: Papers and Proceedings* 2009, 99:2, 198-204.
10. Dell, Melissa, Benjamin F. Jones and Benjamin A. Olken. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics* 2012, 4(3): 66-9
11. Dell, Melissa, Benjamin F. Jones and Benjamin A. Olken.. "What do We Learn from the Weather? The New Climate-Economy Literature?" *Journal of Economic Literature* 2014, 52(3), 740-798.
12. Deryugina, Tatyana, and Solomon M. Hsiang. "Does the Environment Still Matter? Daily Temperature and Income in the United States". No. w20750. National Bureau of Economic Research, 2015.
13. Deschnes, Olivier, and Michael Greenstone. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics* 3, no. 4 (2011): 152-185.
14. Energy Information Agency. (2009) "Residential Energy Consumption Survey: 2009". <http://www.eia.gov/consumption/residential/about.cfm>

15. Graff Zivin and Matthew Neidell. "Environment, Health and Human Capital". *Journal of Economic Literature* 2013, 51(3), 689-730.
16. Graff Zivin, Joshua and Matthew Neidell. "Temperature and the Allocation of Time: Implications for Climate Change", *Journal of Labor Economics*, Vol. 32, No. 1 (January 2014), pp. 1-26
17. Hartwick, J. M. (1977). "Intergenerational equity and the investing of rents from exhaustible resources." *The American Economic Review*: 972-974.
18. Heal, Geoffrey, and Jisung Park. "Reflections: Temperature Stress and the Direct Impact of Climate Change: A Review of an Emerging Literature." *Review of Environmental Economics and Policy* 10.2 (2016): 347-362.
19. Hornbeck, Richard. "The enduring impact of the American Dust Bowl: Short-and long-run adjustments to environmental catastrophe." *The American Economic Review* 102.4 (2012): 1477-1507.
20. Hsiang, S. M. (2010). "Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America." *Proceedings of the National Academy of Sciences*, 107(35), 15367-15372.
21. Kahn, M. E. (2005). "The death toll from natural disasters: the role of income, geography, and institutions." *Review of Economics and Statistics* 87(2): 271-284.
22. Kahn, M.E. *The Climate Change Adaptation Literature*. *Rev Environ Econ Policy* 2016; 10 (1): 166-178. doi: 10.1093/reep/rev023
23. Lobell, David B., and Gregory P. Asner. "Climate and management contributions to recent trends in US agricultural yields." *Science* 299, no. 5609 (2003): 1032-1032.
24. Mackworth, N. (1946). "Effects of heat on wireless operators." *British journal of industrial medicine* 3(3): 143.
25. Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw. "The impact of global warming on agriculture: a Ricardian analysis." *The American economic review* (1994): 753-771.
26. Mendelsohn, R., Morrison, W., Schlesinger, M. E., & Andronova, N. G. (2000). "Country-specific market impacts of climate change". *Climatic change*, 45(3-4), 553-569.
27. Olmstead, A. L. and P. W. Rhode (2011). "Adapting North American wheat production to climatic challenges, 1839?2009." *Proceedings of the National Academy of Sciences* 108(2): 480-485.
28. Park, J., & Heal, G. (2013). "Feeling the heat: Temperature, physiology & the wealth of nations" (No. w19725). National Bureau of Economic Research.
29. Samuelson, Paul Anthony. "Foundations of economic analysis." (1948).
30. Schelling, T. C. (1992). "Some economics of global warming." *The American Economic Review*: 1-14.

31. Schlenker, W. and M. J. Roberts (2009). “Nonlinear temperature effects indicate severe damages to US crop yields under climate change.” *Proceedings of the National Academy of Sciences* 106(37): 15594-15598.
32. Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher. “Will US agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach.” *American Economic Review* (2005): 395-406.
33. Seppanen, O., et al. (2006). “Effect of temperature on task performance in office environment.” Lawrence Berkeley National Laboratory.
34. Sudarshan, Anant, and Meenu Tewari. “The economic impacts of temperature on industrial productivity: Evidence from indian manufacturing”. Working Paper, 2014.

Tables and Figures




Figure 1: Figure 1. Stylized representation of the potential bias in estimating climate damages without taking future adaptation into account, assuming that adaptation can reduce impacts in the longer term.

	Short Run		Long Run	
Secular	Intra-day activity substitution	Daily labor supply (extensive margin)	AC stock	Migration
	Labor effort	Inter-day activity substitution	Clothing purchases	Genetic adaptation
		AC flow	Occupational choice	
			Labor supply extensive (retirement)	
Directed	Workplace regulations	Labor agreements (e.g. contract structure)	R&D in reducing adaptation stock costs	Social Norms
		R&D in reducing adaptation flow cost (energy cost)	Infrastructure	
			Migration assistance	

Table 1: Possible adaptation mechanisms in response to temperature stress

Notes: Some potential adaptation mechanisms, organized along the following dimensions of secular versus directed, short-run versus long-run.

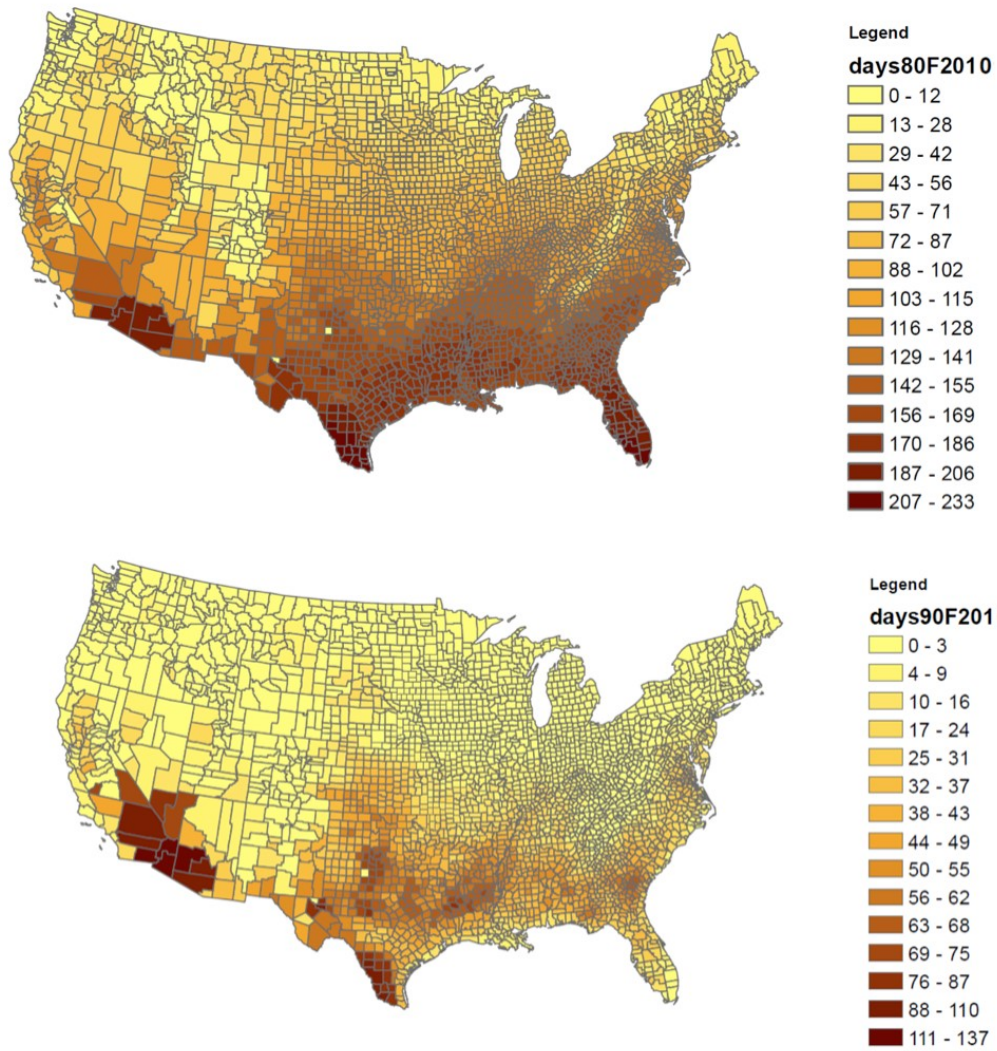


Figure 2: Incidence of Hot Days per Year.

Notes: Top panel shows number of days with daily max temperature above 80°F in 2010.

Bottom panel shows number of days with daily max temperature above 90°F in 2010.

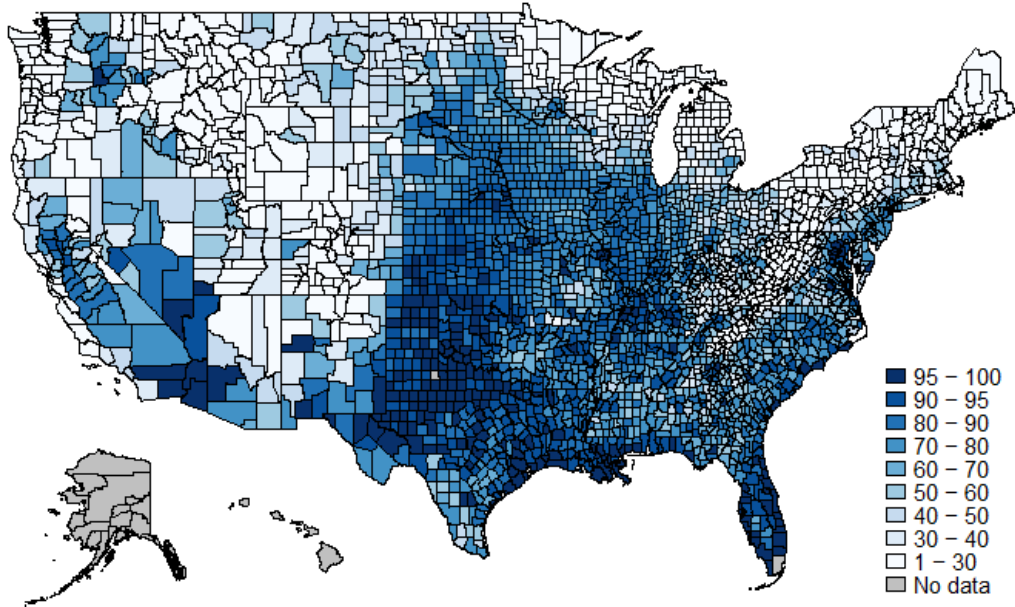


Figure 3: Total Residential AC Penetration (window and central) by county in 1990.
 Notes: Imputed residential AC penetration rates (in percentage of households) by county in 1990. Includes residential window and central AC units.

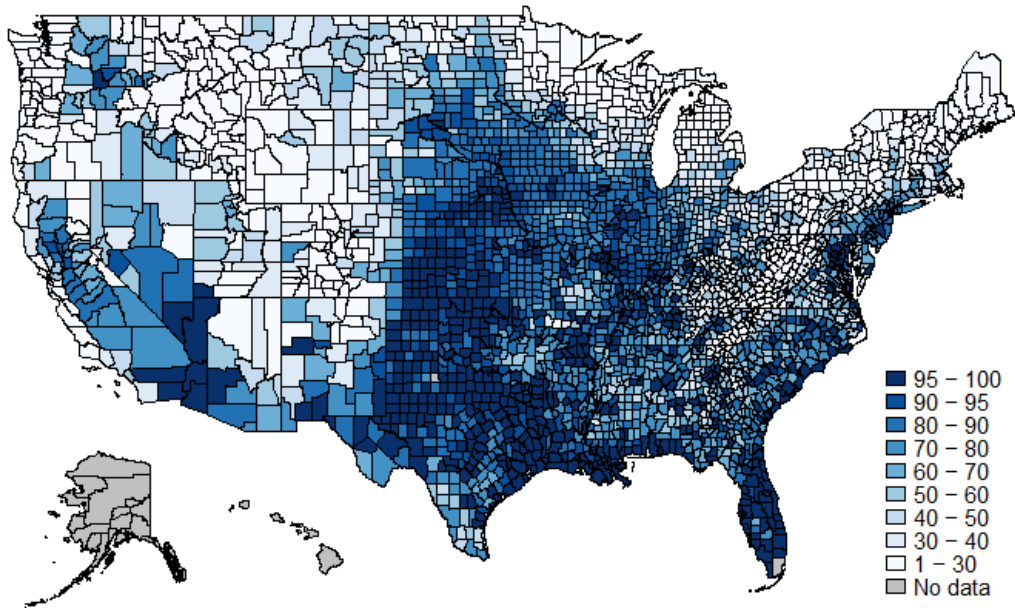


Figure 4: Total Residential AC Penetration (window and central) by county in 2000.
 Notes: Imputed residential AC penetration rates (in percentage of households) by county in 2000. Includes residential window and central AC units.

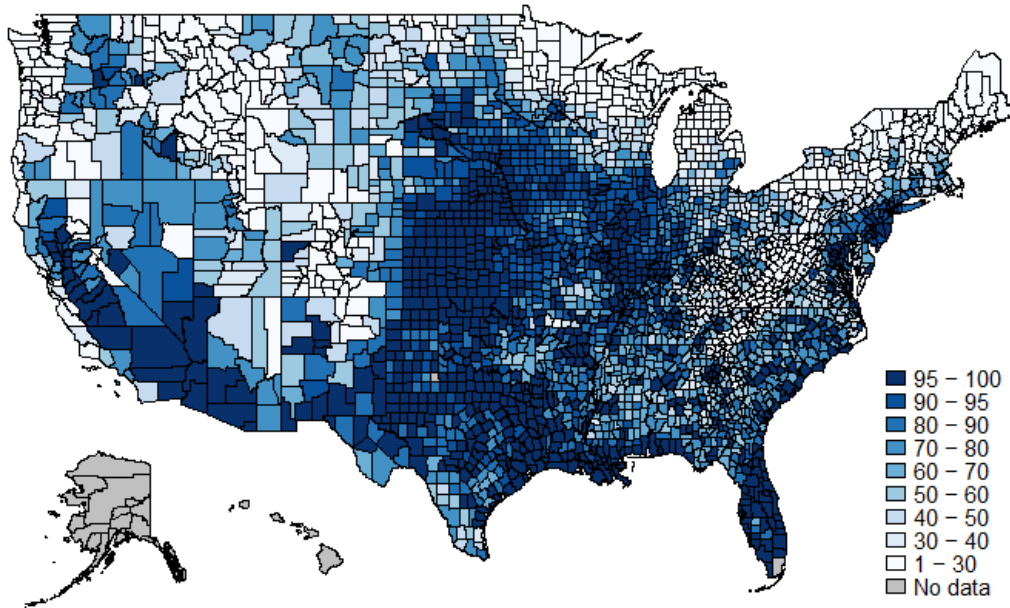


Figure 5: Total Residential AC Penetration (window and central) by county in 2010. Notes: Imputed residential AC penetration rates (in percentage of households) by county in 2010. Includes residential window and central AC units.




Figure 6: Imputed AC penetration in New York County, NY, and Grady County, GA, 1986-2012. Includes window units and central AC for residential dwellings. County-level base values taken from 1980 census. Annual rates of change taken at the census region level from RECS (2012).

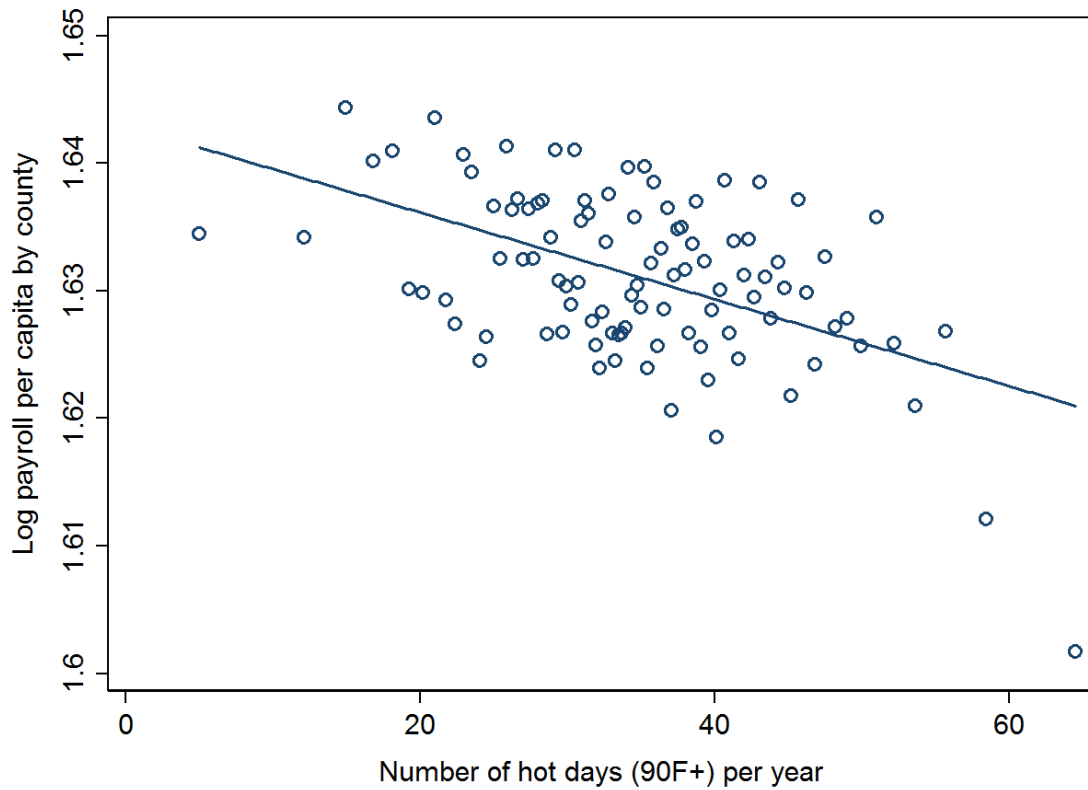


Figure 7: Payroll and Temperature

Notes: Residualized variation including county and year fixed effects, state-specific cubic time trends and non-parametric controls for all other degree days. Days with maximum temperature between 70-79°F is the omitted category.

Table 2: Measuring the Impact of Short-Run Heat Exposure on Annual Production.

	(1)	(2)	(3)	(4)
90°F and above	-0.000264** (0.0000933)	-0.000345*** (0.0000878)	-0.000262** (0.0000931)	-0.000340*** (0.0000918)
0°F to 9°F	0.000140 (0.000618)	0.0000730 (0.000582)	0.000137 (0.000618)	0.0000674 (0.000609)
10°F to 19°F	0.000676 (0.000379)	0.000634 (0.000351)	0.000677 (0.000379)	0.000633 (0.000367)
20°F to 29°F	0.0000126 (0.000249)	0.000230 (0.000249)	0.0000190 (0.000249)	0.000237 (0.000260)
30°F to 39°F	0.0000990 (0.000201)	0.000114 (0.000195)	0.000101 (0.000201)	0.000114 (0.000204)
40°F to 49°F	-0.0000111 (0.000182)	0.000136 (0.000180)	-0.00000849 (0.000182)	0.000139 (0.000189)
50°F to 59°F	0.000182 (0.000150)	0.000154 (0.000142)	0.000184 (0.000150)	0.000155 (0.000148)
60°F to 69°F	0.00000749 (0.000122)	0.0000311 (0.000115)	0.00000904 (0.000122)	0.0000327 (0.000120)
Avg Precip	-0.00313 (0.00245)	-0.00406 (0.00228)	-0.00314 (0.00245)	-0.00404 (0.00238)
Extreme Precip Events	-0.00152 (0.00467)	-0.00340 (0.00389)	-0.00150 (0.00467)	-0.00337 (0.00407)
N	72781	72781	72781	72781
r2	0.993	0.996	0.993	0.996
Fixed Effects				
County	X	X	X	X
Year	X	X	X	X
State-specific linear trends	X			
County-specific linear trends	X	X		
State-specific cubic trends			X	
County-specific cubic trends				X

Robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The dependent variable is log non-agricultural payroll per capita. County and year fixed effects suppressed in output. Robust standard errors are clustered at the state-by-year level. The number of days with temperatures in the 70° to 79° bin is the omitted category.



Figure 8: Payroll and Temperature in Highly Exposed Industries.

Notes: Residualized variation including county and year fixed effects, and state-specific cubic time trends and non-parametric controls for all other degree days. Days with maximum temperature between 70-79°F as omitted category.




Figure 9: Payroll and Temperature in Non-Exposed Industries.

Notes: Residualized variation including county and year fixed effects, and state-specific cubic time trends and non-parametric controls for all other degree days. Days with maximum temperature between 70-79°F as omitted category.




Figure 10: Payroll and Temperature in the Construction Sector

Notes: Residualized variation including county and year fixed effects, and state-specific cubic time trends and non-parametric controls for all other degree days. Days with maximum temperature between 70-79°F as omitted category.

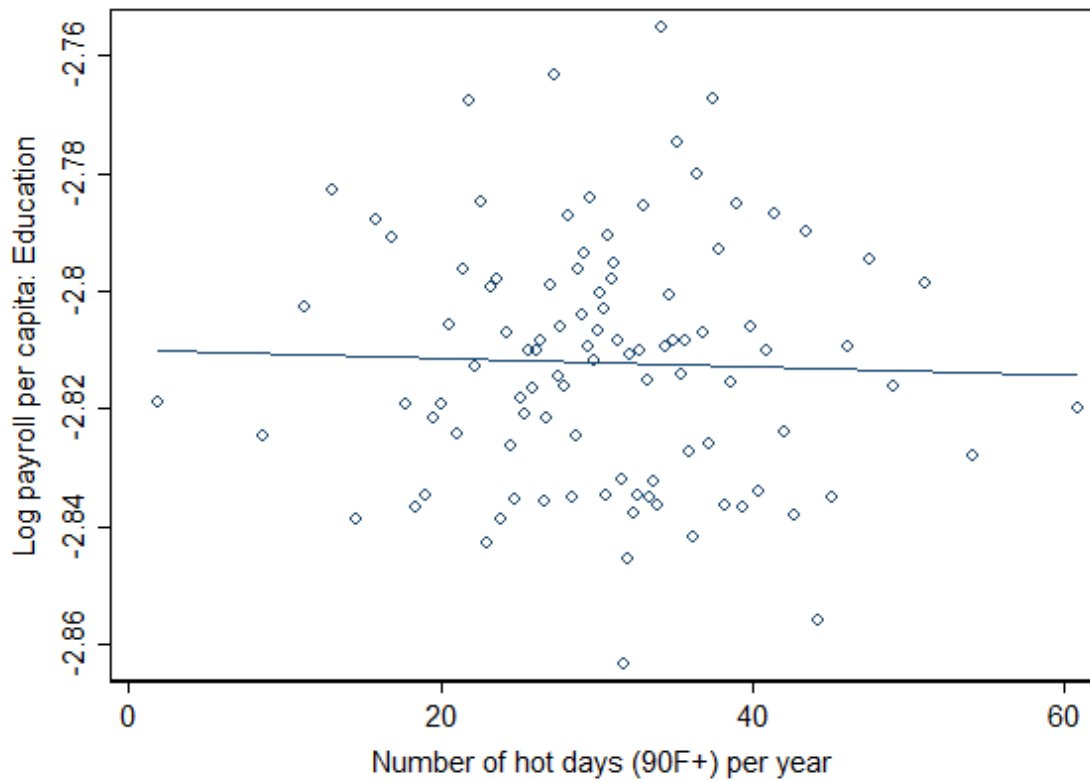


Figure 11: Payroll and Temperature in the Education Sector

Notes: Residualized variation including county and year fixed effects, and state-specific cubic time trends and non-parametric controls for all other degree days. Days with maximum temperature between 70-79°F as omitted category.



Figure 12: Payroll and Temperature in the Healthcare Sector

Notes: Residualized variation including county and year fixed effects, and state-specific cubic time trends and non-parametric controls for all other degree days. Days with maximum temperature between 70-79°F as omitted category.

Table 3: Production Impacts by Quintile of Climate Distribution.

	(1) Quintile 1	(2) Quintile 2	(3) Quintile 3	(4) Quintile 4	(5) Quintile 5
90° and above	-0.00352** (0.00121)	-0.00202** (0.000729)	-0.00168** (0.000517)	-0.00144*** (0.000368)	-0.00138** (0.000534)
0°F to 9°F	-0.000831 (0.00184)	0.00122 (0.00245)	0.0126** (0.00389)	0.00677 (0.00836)	0.0391 (0.0295)
10°F to 19°F	-0.0000766 (0.00113)	0.00116 (0.00149)	0.00236 (0.00220)	-0.00131 (0.00410)	-0.00296 (0.0167)
20°F to 29°F	-0.000523 (0.000816)	-0.00153 (0.000929)	0.00542*** (0.00148)	-0.00200 (0.00180)	-0.00277 (0.00612)
30°F to 39°F	-0.000689 (0.000662)	-0.000453 (0.000719)	0.00266* (0.00107)	0.000505 (0.00114)	0.000136 (0.00271)
40°F to 49°F	-0.000239 (0.000627)	-0.000456 (0.000701)	0.00189* (0.000937)	-0.00171 (0.00101)	-0.00289 (0.00164)
50°F to 59°F	0.000387 (0.000536)	0.000392 (0.000625)	0.00158 (0.000811)	-0.000647 (0.000781)	-0.00257* (0.00102)
60°F to 69°F	0.0000279 (0.000452)	0.0000208 (0.000555)	0.00107 (0.000638)	-0.000142 (0.000696)	-0.000359 (0.000855)
Avg Precip	-0.00259 (0.000950)	-0.0160 (0.0111)	-0.0208 (0.0129)	-0.0201 (0.0106)	-0.0108 (0.0161)
Extreme Precip Events	0.0585** (0.0225)	0.0215 (0.0215)	-0.0124 (0.0258)	0.0508* (0.0217)	-0.0521 (0.0427)
N	14329	14242	13452	13743	13961
r2	0.856	0.869	0.865	0.878	0.855

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The dependent variable is log payroll per capita in highly exposed sectors. All regressions include county and year fixed effects which are suppressed in the output, as well as state-specific cubic time trends in payroll. Robust standard errors are clustered at the state-by-year level. The number of days with temperatures in the 70° to 79° bin is the omitted category.

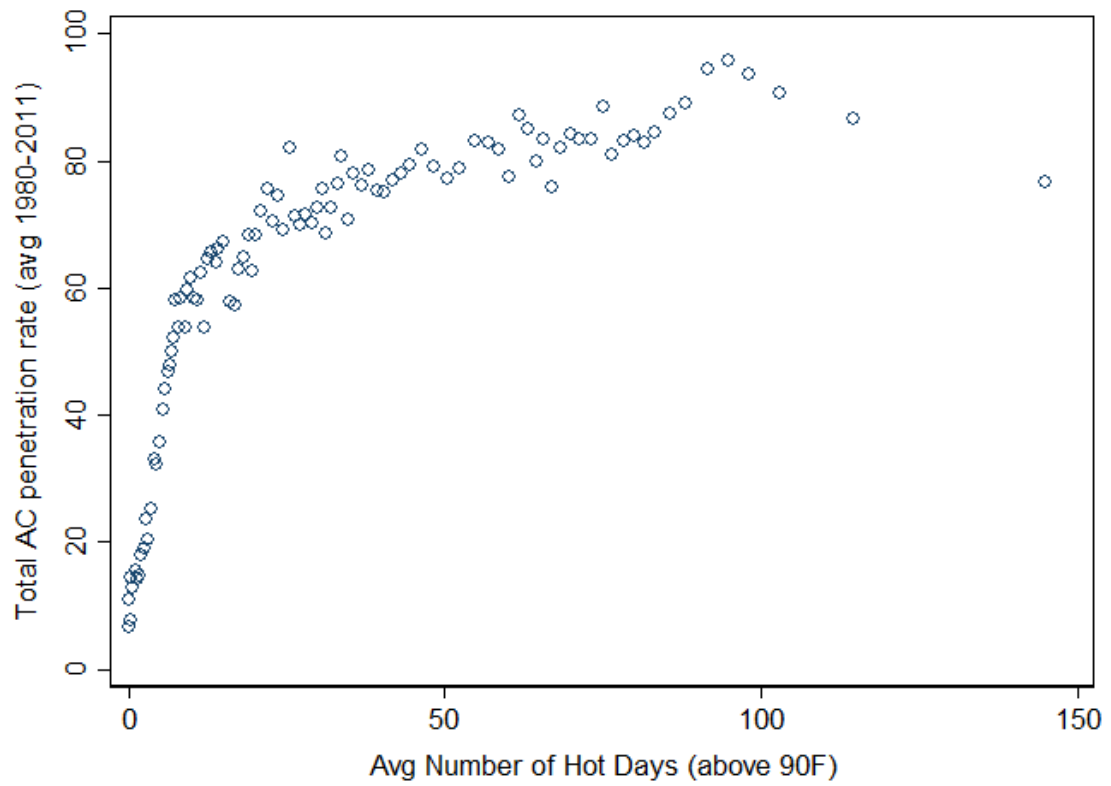


Figure 13: Average AC penetration by county and the average incidence of hot days with temperature above 90°F (1986-2012). Includes controls for per capita income by county.

Table 4: Measuring Adaptation

	(1)	(2)	(3)	(4)	(5)
Above 90°	-0.00222*** (0.000543)	-0.00286*** (0.000781)	-0.00316*** (0.000841)	-0.00281*** (0.000684)	-0.00287*** (0.000686)
Above 90° × Avg Hot Days per Year (90+)	0.0000210* (0.00000988)		0.0000170 (0.0000101)		
Above 90° × Avg Total AC(%)		0.0000231* (0.00000966)	0.0000149 (0.00000922)		
Above 90° × Total AC(%)				0.0000233** (0.00000835)	0.0000242** (0.00000840)
Below 35° × Total AC(%)					0.00000965 (0.00000599)
Total AC(%)				0.00617*** (0.000838)	0.00588*** (0.000875)
N	73219	73219	73219	73219	73219
r ²	0.862	0.862	0.862	0.862	0.862

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: The dependent variable is log payroll per capita in highly exposed sectors. All regressions include county and year fixed effects which are suppressed in the output, as well as state-specific cubic time trends in payroll. Robust standard errors are clustered at the state-by-year level. The number of days with temperatures in the 70° to 79° bin is the omitted category.