NBER WORKING PAPER SERIES

ESTIMATING THE CONSEQUENCES OF CLIMATE CHANGE FROM VARIATION IN WEATHER

Derek Lemoine

Working Paper 25008
http://www.nber.org/papers/w25008

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2018, revised May 2024

I thank numerous seminar participants for valuable comments. I thank Olivier Deschênes, Ryan Kellogg, and Wolfram Schlenker for helpful conversations and Wint Thu for assistance compiling weather and climate model data. An early version circulated as "Sufficient Statistics for the Cost of Climate Change." The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Estimating the Consequences of Climate Change from Variation in Weather
Derek Lemoine
NBER Working Paper No. 25008
September 2018, revised May 2024
JEL No. D84,H43,Q54

ABSTRACT

I formally define the limits of what we can learn about the consequences of long-run climate change from short-run weather shocks. I show that conventional approaches to estimating climate impacts require assuming that payoffs are independent of capital and resource stocks. I derive a new indirect least squares estimator that bounds long-run climate impacts from short-run responses to weather under less restrictive assumptions. In an application to the U.S. economy, I project that climate change will reduce steady-state income per capita by at least 1.8% in the Midwest, by at least 1% in the Northeast, and by at least 0.23% in the West. Each lower bound implies damages beyond the 95% confidence intervals produced by the conventional approach.

Derek Lemoine
Department of Economics
University of Arizona
McClelland Hall 401EE
Tucson, AZ 85721
and CEPR
and also NBER
dlemoine@email.arizona.edu

A data appendix is available at http://www.nber.org/data-appendix/w25008

1 Introduction

Ignorance of the economic costs of climate change prevents economists from giving firm policy guidance. A pressing research agenda seeks to estimate these costs. Recognizing that different locations have different climates, many economists have hoped to estimate the effects of climate change from the correlation between climate and outcomes of interest over space (e.g., Mendelsohn et al., 1994; Schlenker et al., 2005; Nordhaus, 2006). However, locations differ in many ways, making it difficult to isolate the effects of being in one climate or another.[1]

Intriguingly, though, the same location does experience different weather at different times. Stimulated by Deschênes and Greenstone (2007) and Dell et al. (2012), a rapidly growing empirical literature estimates the consequences of a location happening to experience cooler-than-average or hotter-than-average weather.[2] Researchers project the consequences of climate change by combining credibly estimated effects of weather from panel regressions with scientists' predictions about how climate change will alter the distribution of weather. But in the absence of a formal model, it has been unclear what assumptions are required in order to justify the extrapolation of weather impacts to climate change impacts.[3]

This paper formally relates the effects of climate change to the effects of weather shocks. I focus on the dynamic structures of weather shocks and climate change: weather shocks are transient, whereas climate change permanently alters the distribution of weather. I show that estimating long-run effects of climate change requires estimating the direct effects of altered average weather and the average effects of adapting to altered weather, which encompasses both ex-post adaptation (through which agents react to altered weather realizations) and ex-ante adaptation (through which agents anticipate the altered distribution of future weather).

Within this model, I formally derive the effects of climate change that empirical researchers aim to estimate and express panel regressions' coefficients in terms of model primitives. I show that conventional methods require strict assumptions if their calculations are to be relevant for evaluating climate change. In particular, these methods require that economic decisions are independent over time, as when there are no capital or resource stocks. Relaxing this assumption, I show that the best possible weather regressions differ from the ones that empirical researchers typically run by including lags of weather and including forecasts of weather. However, I also

[1]See Dell et al. (2014) and Auffhammer (2018) for expositions and Massetti and Mendelsohn (2018) for a review.

[2]For reviews, see Dell et al. (2014), Carleton and Hsiang (2016), Heal and Park (2016), and Carter et al. (2018). Blanc and Schlenker (2017) and Kolstad and Moore (2020) discuss the strengths and weaknesses of relying on panel variation in weather.

[3]For instance, Dell et al. (2014, 771–772) emphasize that "short-run changes over annual or other relatively brief periods are not necessarily analogous to the long-run changes in average weather patterns that may occur with climate change." And Mendelsohn (2019, 272) observes, "An important failing of current weather panel studies is that they lack a clear theoretical model."

show that even the best possible weather regressions suffer from three biases when used to project climate change impacts.

All three biases derive from the adaptation channel. These biases reflect the possibilities that (i) agents may not pay the fixed costs of modifying long-lived infrastructure in response to a transient weather shock but would pay those costs in response to a permanent change in climate; (ii) a shock to short-run forecasts is not completely equivalent to altering expectations of weather in every period; and, of most interest, (iii) reactions to short-run weather shocks are not identical to reactions to living with altered weather period after period. Actions are intertemporal complements (substitutes) if undertaking more actions in one period increases (decreases) optimal actions in later periods through a stock variable. For example, actions are intertemporal complements when they represent capital investment in the presence of adjustment costs and are intertemporal substitutes when they deplete a scarce resource stock. Channel (iii) depends on which case holds. In the former case, estimates derived from short-run weather variation can understate long-run adaptation to climate change because agents have more flexibility in the long run, but in the latter case, estimates derived from short-run weather variation overstate long-run adaptation to climate change because agents have a hard time maintaining adaptation responses.[4]

What, then, is an empirical researcher to do? I develop a new indirect least squares estimator of climate impacts (Tinbergen, 1930, 1995) that substantially relaxes the assumptions under which panel variation can identify effects of climate change. I show that the estimator can partially identify the long-run effects of climate change from short-run weather shocks, even without observing the actions agents and firms could choose, without observing the capital or resource stocks that they interact with, and without parameterizing payoffs or transition equations for stock variables. I derive this estimator in four steps. First, I express climate change impacts in terms of my setting's theoretical primitives. Second, I express reduced-form weather coefficients in terms of the same theoretical primitives. This system is just-identified in that there are as many aggregates of unknown theoretical primitives as there are equations, and it so happens that these same aggregates of theoretical primitives appear in the climate change expression from the first step. So, third, I solve for these aggregates of theoretical primitives in terms of the estimable reduced-form coefficients. Finally, I use this solution to calculate climate impacts from the analytical expression derived in the first step and the estimable reduced-form weather coefficients.

The identification is purely reduced-form, as the only point of contact with data is a fixed effects regression that relies on conventional panel variation in weather. The-

[4]Both types of stories exist in the literature (see Auffhammer, 2018). For instance, in studies of the agricultural impacts of climate change, Deschênes and Greenstone (2007) conjecture that long-run adjustments to changes in climate should be greater than short-run adjustments to weather shocks because there may be costs to adjusting crops, whereas Fisher et al. (2012) and Blanc and Schlenker (2017) conjecture that constraints on storage and groundwater pumping, respectively, could make short-run adjustments exceed long-run adjustments.

ory determines both the specification of the regression and the calculations that use its estimates.[5] Intuitively, variation in contemporary weather identifies a combination of direct weather effects and ex-post adaptation. Past weather identifies ex-post adaptation because it affects current payoffs only through the effects of ex-post adaptation on capital or resource stocks. I identify the direct effects of weather by combining the coefficient on contemporary weather with the estimate of ex-post adaptation. Forecasts identify ex-ante adaptation because they affect current payoffs only when agents take actions based on forecasts.[6] And comparing the effects of weather over shorter and longer lags identifies whether estimates derived from short-run weather variation tend to over- or underestimate the effects of longer-run adaptation.

I demonstrate the new method using a panel analysis of the effect of temperature changes on changes in U.S. counties' output per capita from 2002–2019 and on changes in U.S. counties' income per capita from 1970–2019.[7] Wald tests reject the hypothesis that the effects of lagged extreme heat are jointly zero in the data. As a result, conventional calculations would not be informative about climate change impacts in this application.

Applying my new indirect least squares estimator, I estimate that extreme heat significantly reduces income per capita in the Midwest, Northeast, and West and significantly increases income per capita in the South. Combining all weather variables, my central estimates suggest that end-of-century climate change will reduce output (income) per capita by 4% (2%) in the Midwest and 0.5% (2%) in the Northeast. Projected effects on output per capita are small and of ambiguous sign in the South and West (possibly reflecting the short panel), but central estimates for income per capita project benefits of 0.3% in the South and losses of 0.6% in the West.

In most regions, the theoretically-grounded indirect least squares estimator predicts effects on income per capita that are significantly different from what conventional methods would predict. I find that the conventional approach underestimates losses from extreme heat in most regions because it relies on an estimated coefficient that entangles direct costs of temperature shocks with short-run benefits from ex-post adaptation. For optimizing agents, those short-run benefits come at the expense of long-run costs that are identified by the coefficient on lagged weather. My the-

[5] Critically, this calculation does not require the specification of structural parameters. This approach is in the spirit of Marschak's Maxim. Heckman (2010, 359) writes, "Marschak's Maxim suggests that economists should solve well-posed economic problems with minimal assumptions. All that is required to conduct many policy analyses or to answer many well-posed economic questions are policy invariant combinations of the structural parameters that are often much easier to identify than the individual parameters themselves and that do not require knowledge of individual structural parameters." It is also related to sufficient statistics approaches (see Chetty, 2009) and to price theory (see Weyl, 2019).

[6] Shrader (2020) also uses variation in forecasts to identify ex-ante adaptation, in an analysis not focused on climate change.

[7] Colacito et al. (2019) estimate negative effects of summer temperature on output growth in U.S. states. Barker (2022) critique their methods and interpretation. In a working paper, Deryugina and Hsiang (2017) estimate effects on county income per capita over 1969–2011.

oretical analysis shows how to combine the estimated effects of contemporary and lagged weather to isolate both the direct effects of temperature changes and the long-run costs of ex-post adaptation. The isolated direct effects drive my estimated total effects.

A ratio of lagged effects of weather determines whether estimated costs of ex-post adaptation are upper or lower bounds on steady-state costs (i.e., whether actions are intertemporal substitutes or complements). Estimates of this ratio can be imprecise when the longer lags have ambiguous sign, which is a particular problem in the short panel available for output per capita. In the longer panel available for income per capita, estimates in the longer panel for income per capita are consistent with actions being intertemporal complements in the Northeast and South, as under an adjustment cost model. In a supplemental industry-level analysis, I report evidence consistent with actions being intertemporal complements in the two industries (agriculture and retail) that show evidence of temperature impacts. In such cases, steady-state ex-post adaptation is greater than estimated here, so that estimated losses are lower bounds on steady-state effects. From the smaller end of the 95% confidence interval, the ILS estimator projects losses of at least 1.8% in the Midwest, at least 1% in the Northeast, and at least 0.23% in the West. These lower bounds suggest worse damages than even the 95% confidence intervals produced by conventional methods. Using the new estimator matters both for interpretability in terms of climate change and for the numbers produced.

Despite the importance of empirically estimating the costs of climate change and the sharpness of informal debates around the relevance of the burgeoning weather regression literature to climate change, there has been remarkably little prior formal analysis of the economic distinction between weather shocks and climate change.[8] The exceptions are Hsiang (2016) and Deryugina and Hsiang (2017). They argue that the simplest weather regression exactly identifies the effect of climate on payoffs. Their setting assumes that choices made in one period are completely independent of choices made in any other period and that outcomes and actions depend only on the distribution of weather (i.e., only on the climate), not on the weather realized from this distribution. Their formal setting is a special case of the present setting (see footnote 20 in Section 3). I show that their optimistic result does not survive generalizing payoffs to depend on a stock variable, such as capital or natural resources. Whereas the envelope theorem implies that changes in actions have no effects on payoffs in a static environment like that of Hsiang (2016) and Deryugina and Hsiang (2017), the envelope theorem in a dynamic environment applies to the intertemporal value function, not to the flow payoffs that empirical researchers typically observe.

Motivated by concerns that responses to weather shocks may not reflect responses

[8] Mérel and Gammans (2021) explore the conditions under which actions chosen under the full distribution of weather identify actions chosen for average weather, assuming that the envelope theorem prevents actions from having first-order consequences for the cost of climate change. Carter et al. (2018) discuss several econometric issues in the estimation of panel models of weather.

to climate change, recent literature takes a variety of approaches to avoid relying exclusively on short-run variation in weather. First, recent reduced-form work estimates how the marginal effect of weather varies cross-sectionally with a location's climate (e.g., Auffhammer, 2022; Carleton et al., 2022). This approach forsakes the clean identification of panel variation. I here explore the limit of what researchers can learn from purely panel variation in weather so that identification is never in question. Second, some reduced-form work uses "long difference" estimators that aggregate weather over larger timesteps (e.g., Dell et al., 2012; Burke and Emerick, 2016). Appendix A shows that long difference estimators inherit the biases suffered by standard weather regressions (see footnote 31). Third, recent work aggregates lags of weather in an attempt to obtain time series variation in a location's climate (e.g., Bento et al., 2023; Leduc and Wilson, 2023; Mohaddes et al., 2023). I here show how variation in shorter lags of weather, which is more likely to be exogenous and observable, is critical to theoretically-grounded extrapolation from weather to climate. Finally, recent work accounts for adaptation by specifying and calibrating macroeconomic models (e.g., Fried, 2022; Bakkensen and Barrage, 2024). I impose less structure on the economic environment and explore the limits of reduced-form estimators that retain the quasi-experimental variation of panel models.

The challenge of attempting to estimate long-run effects from short-run variation is a common one in empirical economics. For instance, labor economists desire the long-run consequences of changing the minimum wage, but inflation converts observed minimum wage increases into short-run shocks (Sorkin, 2015).[9] And macroeconomists formerly hoped to learn about long-run output-inflation tradeoffs by estimating distributed lag models, but Lucas (1972) argued that, when agents have rational expectations, the response to a transient inflation shock is not informative about the long-run effects of permanently changing inflation policy. Here we desire the long-run effect of changing the "policy rule" used by nature to generate weather.

The next section describes the setting. Section 3 analyzes a special case without dynamic linkages, as in prior literature. Section 4 analyzes the full model and delineates what we can learn from reduced-form regressions. Section 5 derives the indirect least squares estimator. Section 6 applies the new estimator to U.S. counties' output and income per capita. The final section outlines potential extensions. The appendix contains supplemental results and proofs.

[9] Three other papers are related to both Sorkin (2015) and the present paper's project. First, I here formalize analogues to arguments in Hamermesh (1995) about why the pre- and post-periods around a minimum wage increase are not true pre- and post-periods. Second, in a model of dynamic stock accumulation, Hennessy and Strebulaev (2020) show that estimated responses to transient shocks can differ substantially from the theory-implied causal effects that empirical researchers seek to test. The present paper is similar in deriving sufficient conditions for estimated effects to match theory-implied effects. Third, Keane and Wolpin (2002) describe tradeoffs between cross-sectional and panel variation when estimating the effects of welfare benefits. These tradeoffs are similar to those that motivate the present paper.

2 Setting

In each period t, agents receive payoffs $\pi(w_t, A_t, S_t; K)$, with π quadratic in each of the four arguments.[10] After observing weather w_t, agents choose actions A_t as a form of adaptation, where $\pi_{AA} < 0$ (subscripts indicate partial derivatives). Agents can also affect a stock variable S_t, where $\pi_S = 0$ in Section 3 and $\pi_{SS} < 0$ everywhere else. The stock evolves as $S_{t+1} = g\, S_t + h(A_t)$, with h monotonic.[11] The parameter $g \in [0, 1)$ controls the persistence of actions. If $g = 0$, the time $t + 1$ stock depends only on time t actions, as with acreage planted versus fallowed. If $g > 0$, the time $t + 1$ stock depends on all past actions, as with a capital stock that depends on past investments and overdrafts from a groundwater reservoir that recharges towards its steady state.

The stock can affect an agent's payoffs from pursuing different actions. When $h'\,\pi_{AS} < 0$, actions are ***intertemporal substitutes***, so that choosing a higher action in one period reduces the marginal benefit of actions in the subsequent period. I describe this case as a resource scarcity story. For instance, pumping groundwater today raises the cost of pumping groundwater tomorrow. When $h'\,\pi_{AS} > 0$, actions are ***intertemporal complements***, so that choosing a higher action in one period increases the marginal benefit of actions in the subsequent period. I describe this case as an adjustment cost story because it favors approaching a high action via a sequence of smaller steps. For instance, small changes to capital stocks may be easier to implement than large changes. The magnitude of $h'\,\pi_{AS}$ affects how agents prepare in advance of a weather event that they know will change their preferred actions. When $|h'\,\pi_{AS}|$ is large, agents prefer to begin adapting actions before a weather event arrives, but when $|h'\,\pi_{AS}|$ is small, agents may wait to undertake most adaptation only once a weather event has arrived.

Agents understand the climate C, which controls the distribution of weather. At all times before $t - 2$, an agent's only information about time t weather consists in knowledge of the climate. However, at time $t - 2$ specialized information about time t weather becomes available in the form of a random variable $\epsilon_{2,t-2}$. The agent uses this information to form a forecast $f_{2,t-2}$ of time t weather: $f_{2,t-2} = C + \zeta\epsilon_{2,t-2}$.[12] The parameter $\zeta \geq 0$ is a perturbation parameter that will be useful for analysis (see Judd, 1996). At time $t - 1$, the agent receives additional news about time t weather in the form of a random variable $\epsilon_{1,t-1}$. The agent refines her forecast of time

[10] I refer to "agents" and "actions", but one can also think of firms choosing quantities, with weather affecting profits through prices and/or the production function.

[11] I abstract from externalities in use of the stock and from the possibility that the stock is directly vulnerable to weather shocks. Future work could consider common pool resources and weather-exposed stocks.

[12] Implicitly, $f_{k,t} = C$ for $k > 2$. Results generalize straightforwardly when extending the analysis to allow for specialized forecasts of weather more than two periods away. Because doing so generates little new insight but imposes additional notation, I restrict attention to the case with specialized forecasts beginning only two periods ahead of a weather event.

t weather to $f_{1,t-1} = f_{2,t-2} + \zeta\epsilon_{1,t-1}$. Finally, the agent may be surprised by a random component $\epsilon_{0,t}$ of time t weather, where $w_t = f_{1,t-1} + \zeta\epsilon_{0,t}$. Let the $\{\epsilon_{0,t}, \epsilon_{1,t}, \epsilon_{2,t}\}$ be jointly normal with covariance matrix Σ. Reflecting rationality of beliefs, the $\epsilon_{i,t}$ are mean-zero and uncorrelated with the $\epsilon_{i,t-1}$. Even though the news represented by $\epsilon_{i,t}$ is serially uncorrelated, the weather realizations w_t are serially correlated if Σ is not diagonal.

Each agent chooses actions to maximize the expected present value of payoffs over an infinite horizon:

$$\max_{\{A_t(S_t,w_t,f_{1,t},f_{2,t})\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t E_0\left[\pi(w_t, A_t, S_t; K)\right],$$

where $\beta \in (0,1)$ is the per-period discount factor, E_0 denotes expectations at the time 0 information set, and S_0, w_0, $f_{1,0}$, and $f_{2,0}$ are given. The solution satisfies the following Bellman equation:

$$V(S_t, w_t, f_{1,t}, f_{2,t}; \zeta, K) = \max_{A_t}\left\{\pi(w_t, A_t, S_t; K) + \beta E_t\left[V(S_{t+1}, w_{t+1}, f_{1,t+1}, f_{2,t+1}; \zeta, K)\right]\right\}, \tag{1}$$

where V is the intertemporal value function. With quadratic payoffs, an agent's optimal policy rule will be independent of the covariance matrix Σ, but Σ will affect the agent's optimal policy choices by affecting realizations of weather.[13]

Agents also choose long-lived infrastructure K. This infrastructure represents capital-intensive adaptation that takes years to construct, such as irrigation canals or sea walls (see Aldy and Zeckhauser, 2020). Agents cannot adapt this infrastructure to short-run weather outcomes or forecasts.[14] Infrastructure is chosen before specialized forecasts are available, as its time to build exceeds the horizon of forecasts. Formally, the agent chooses K in a pre-period to solve the following problem:

$$\max_{K}\left\{\sum_{t=0}^{\infty} \beta^t E[\pi(w_t, A_t^*, S_t; K)]\right\},$$

where expectations account for knowledge of the climate but not for specialized forecasts. Together, the decision variables A_t and K bracket the many types of actions

[13]Consistent with much previous literature, climate here controls average weather. One might wonder about the dependence of higher moments of the weather distribution on climate. In fact, the effects of climate change on the variance of the weather are poorly understood and likely to be spatially heterogeneous (e.g., Huntingford et al., 2013; Lemoine and Kapnick, 2016). Further, for economic analysis, we need to know not just how climate change affects the variance of realized weather but how it affects the forecastability of weather: the variance of the weather more than two periods ahead is $\zeta^2 trace(\Sigma)$, so we need to apportion any change in variance between the diagonal elements of Σ (i.e., between each of the $\epsilon_{i,t}$). I leave such an extension to future work.

[14]Such infrastructure is the only kind of action analyzed in previous work that formally relates climate change to weather variation (Hsiang, 2016; Deryugina and Hsiang, 2017).

actual agents may take: actions A_t encompass those implemented immediately and, via $h(\cdot)$, those implemented with a one-period lag based on information about both near-term weather and long-term climate, whereas actions K have a long time to build that makes them insensitive to information about near-term weather.

The setting is meant to be fairly general. The stock can represent non-physical stocks such as household wealth, and actions can represent allocation problems within an economy that is constrained-efficient, as when it solves the problem of a social planner who cannot address externalities (see Section 6). To fix ideas, consider a few examples pertinent to previous empirical research.

In an agricultural application, weather could affect yields or crop prices, actions could be irrigation or planting decisions, the stock could be water supplies or machinery, and long-lived infrastructure could be irrigation canals or available crop varietals.[15] In a flooding application, weather could affect rental value, actions could be investments in the property, the stock could be the quality of the property, and long-lived infrastructure could be sea walls. In a migration application, weather could affect wages, actions could be the choice of future location, and the "stock" (i.e., the inherited state) could be one's current location.[16] In a health application, weather could affect mortality, individuals could organize their outdoor activities around weather forecasts in order to maximize utility net of mortality risks, and they could find their flexibility restricted as more days go by and the stock of postponed activities accumulates (as in Graff Zivin and Neidell, 2009). In an innovation application, weather could affect the value of patents, the action could be investing in research, and the stock could be existing patents on adaptation technologies. In a labor or energy application, weather could affect labor productivity or residential comfort, the action could be using air conditioning, and the stock could be air conditioning equipment. And in a macro application, weather could affect profits via prices, actions could be factor use, and the stock could be capital.

I assume the following technical conditions in settings with $\pi_S \neq 0$. The first ensures that the payoff function is strictly concave in S_t and S_{t+1}, which in turn ensures that there is a uniquely optimal action (Appendix F.3):

$$[\pi_{AS}]^2 < \pi_{SS} \left[\pi_{AA} - \frac{h''(A_t)}{h'(A_t)} \pi_A \right]. \tag{2}$$

Inequality (2) and $\pi_{SS} < 0$ imply

$$\pi_{AA} - \frac{h''(A_t)}{h'(A_t)} \pi_A < 0, \tag{3}$$

[15] Recent literature reports that actions such as irrigation choices and crop substitution (Cui, 2020), acreage planted (Aragón et al., 2021), and pesticide use and weeding effort (Jagnani et al., 2021) respond to weather.

[16] The choice of future location fits the formal framework if it depends on the current location's present and forecasted weather and on the alternate location's average weather.

which is satisfied when h'' is small. The next two conditions ensure that a steady state exists in a deterministic system with $\zeta = 0$ (Appendix F.4):

$$\lim_{A_t \to -\infty} \left[(1-\beta g)\pi_A\left(C, A_t, \frac{h(A_t)}{1-g}; K\right) + \beta h'(A_t)\pi_S\left(C, A_t, \frac{h(A_t)}{1-g}; K\right) \right] > 0, \quad (4)$$

$$\lim_{A_t \to \infty} \left[(1-\beta g)\pi_A\left(C, A_t, \frac{h(A_t)}{1-g}; K\right) + \beta h'(A_t)\pi_S\left(C, A_t, \frac{h(A_t)}{1-g}; K\right) \right] < 0. \quad (5)$$

If $h(\cdot)$ is linear and π_{AS} is zero, these two conditions reduce to $\lim_{A_t \to -\infty} \pi_A > 0$ and $\lim_{A_t \to \infty} \pi_A < 0$. The final condition ensures that the expression for expected optimal actions converges (Lemma 2 in Appendix F.6):

$$h'(A_t)\pi_{AS} \in \Bigg(-\frac{[1 + 2g(1+\beta) + 3\beta g^2]\left[-\pi_{AA} + \frac{h''(A_t)}{h'(A_t)}\pi_A\right] - \beta[h'(A_t)]^2\pi_{SS}}{1+\beta+2\beta g}, \\ \frac{[1 - 2g(1+\beta) + 3\beta g^2]\left[-\pi_{AA} + \frac{h''(A_t)}{h'(A_t)}\pi_A\right] - \beta[h'(A_t)]^2\pi_{SS}}{1+\beta-2\beta g} \Bigg). \quad (6)$$

The interval includes zero. This condition therefore permits both intertemporal complementarity and intertemporal substitutability but limits the degree of either. In sum, all four technical conditions tend to hold when h'' and π_{AS} are not too large.

I am interested in empirical researchers' ability to estimate the consequences of altering C from observable responses to panel variation in w_t and, where available, to panel variation in $f_{1,t}$ and $f_{2,t}$. I assume that empirical researchers observe payoffs, weather, and, possibly, forecasts for J agents (equivalently, firms) in each of T periods. Index these agents by j. To highlight the issue at hand, they are in the same climate C with the same payoff function π but their own stocks. I do not explicitly model the unobservable characteristics that motivate fixed effects specifications because I do not here study whether regression coefficients consistently estimate weather impacts. Instead, I study whether climate change impacts are in principle recoverable from consistently estimated weather impacts.

Working within a formal framework forces us to define the treatment effect of interest. I study the average effects (over time, and thus over weather shocks) of moving agents from one climate to another once agents have had time to adapt to the new climate. This adaptation is based both on experiencing weather drawn from the new distribution of weather and on understanding the distribution of future weather. This climate change treatment is consistent with the dominant exercise in the empirical literature to date, which calculates the effect of replacing a recent distribution of weather with a distribution projected to hold around 100 years later. Following this literature, I will not study how the transition from one climate to another interacts with agents' decisions,[17] or how expectations of a future change

[17]Kelly et al. (2005) frame the cost of learning as an adjustment cost. Quiggin and Horowitz (1999,

in climate affect agents today.[18] These are both important questions but are beyond the scope of the present analysis—and thus far largely beyond the empirical literature that this analysis seeks to inform.

3 Estimating Climate Impacts When There Are No Dynamic Linkages

Begin by considering a setting in which payoffs are independent of the stock S_t: $\pi_S = 0$. Each period's decision problem simplifies to a static problem, with optimal actions $A_t^*(w_t; K)$ satisfying the first-order condition $\pi_A(w_t, A_t^*, \cdot; K) = 0$ and independent of all other periods' actions.[19] The first-order condition for the infrastructure choice problem is $\pi_K(C, A_t^*(C; K), \cdot; K) = 0$.

Define $\bar{A} \triangleq A_t^*(C; K)$ and $\bar{\pi} \triangleq \pi(C, \bar{A}, \cdot; K)$. Appendix F.1 shows that

$$\frac{\mathrm{d}E_0[\pi_t]}{\mathrm{d}C} = \bar{\pi}_w + \underbrace{\bar{\pi}_A \frac{\mathrm{d}\bar{A}}{\mathrm{d}C}}_{=0} + \underbrace{\bar{\pi}_K \frac{\mathrm{d}K}{\mathrm{d}C}}_{=0} = \bar{\pi}_w \qquad \forall t > 2. \tag{7}$$

When agents optimize, the effects of climate on short-run and long-run actions vanish and we need to recover only the direct effect of weather. This envelope theorem intuition is familiar from previous literature (Hsiang, 2016; Deryugina and Hsiang, 2017).[20,21]

Consider the following regression

$$\Delta\pi_{jt} = \alpha_j + \theta \Delta w_{jt} + \eta_{jt}, \tag{8}$$

where $\Delta\pi_{jt} \triangleq \pi_{jt} - \pi_{j(t-1)}$ (and analogously for other variables), α_j is a fixed effect for unit j, and η_{jt} is an error term. Use a hat to denote each estimator. Here and below,

2003) discuss broader costs of adjusting to a change in climate. These papers' adjustment costs are conceptually distinct from the adjustment costs studied here. The present use of "adjustment costs" follows much other economics literature in referring to the cost of changing decisions from their previous levels. I study how these adjustment costs hinder estimation of the consequences of climate change from weather impacts, not how they affect the cost of transitioning from one climate to another.

[18]Severen et al. (2018) show that land markets capitalize expectations of future climate change and correct cross-sectional analyses in the tradition of Mendelsohn et al. (1994) for this effect. I here study responses to widely available, shorter-run forecasts in a longitudinal context and show how to use them to improve panel analyses in the tradition of Deschênes and Greenstone (2007).

[19]Using terminology defined below, ex-ante adaptation is here impossible and ex-post adaptation is here independent of all prior ex-post adaptation.

[20]To recover the setting of Hsiang (2016) and Deryugina and Hsiang (2017), we would have to eliminate the choice of A_t and make π depend on C directly rather than on w_t. In that case, the only available action (the choice of K) is made independently of weather realizations and there is no scope for either ex-post or ex-ante adaptation.

[21]Guo and Costello (2013) show that this envelope theorem intuition breaks down when choice variables are discrete, which could be especially relevant to long-lived infrastructure.

I write regressions in terms of differences for consistency with empirical applications concerned with spurious trends, such as Newell et al. (2021) and Section 6 below.[22] I study the probability limits of this and other estimators as we increase the number of units in the sample. By standard results,

$$\operatorname{plim}\hat{\theta} = \frac{Cov[\Delta\pi_{jt} - \mu_j^\pi, \Delta w_{jt} - \mu_j^w]}{Var[\Delta w_{jt} - \mu_j^w]}, \tag{9}$$

where μ_j^x indicates the mean of variable Δx across time periods within unit j. The following proposition relates this estimator to theoretical primitives.

Proposition 1 (Conventional Estimator)**.** *If* π *is independent of* S_t*, then* $\operatorname{plim}\hat{\theta} = \bar{\pi}_w$.

Proof. See Appendix F.2. Sketch: Expands π to second-order around the deterministic steady state and analyzes (9). □

Therefore, from equation (7),

$$\frac{\mathrm{d}E_0[\pi_t]}{\mathrm{d}C} = \operatorname{plim}\hat{\theta} \qquad \forall t > 2.$$

When payoffs are independent of the stock, the simplest weather regression recovers the average marginal effect of weather and thus recovers the long-run effects of climate.[23]

This is an optimistic result, but this environment with $\pi_S = 0$ is rather specialized. First, we have assumed that history does not matter. Yet capital stocks and storage may adjust only slowly over time and resource constraints may compound over time, as several authors have informally noted (e.g., Deschênes and Greenstone, 2007; Fisher et al., 2012). Capital stocks and resource constraints are intuitively important in many applications, whether agricultural, industrial, or household. Second, we have assumed away any ability to proactively protect oneself against future weather (i.e., to undertake ex-ante adaptation). Yet evidence suggests that farmers adjust planting decisions based on beliefs about the coming season's weather (Rosenzweig and Udry, 2013), fishers adjust plans based on multi-month forecasts of El Niño events (Shrader, 2020), markets price in multi-day hurricane forecasts (Kruttli et al., 2019) and multi-month seasonal climate forecasts (Lemoine and Kapnick, 2024), and people use daily weather forecasts to reduce mortality risk (Shrader et al., 2023). We next turn to the full setting to see how far the optimism engendered by the present specialization has to run.

[22] With one small caveat (see footnote 27 below), the theoretical analysis is unaffected by estimating equations in levels rather than in differences.

[23] Much literature regresses outcomes other than payoffs on weather. In the restricted setting of Section 3, the coefficient on weather in a regression with actions as the dependent variable also recovers the long-run effect of climate on actions.

4 Estimating Climate Impacts in the Presence of Dynamics

In a general environment with $\pi_S \neq 0$, agents' actions depend on realized past weather and on expectations of future weather. Appendix F.4 establishes that the deterministic special case (with $\zeta = 0$ and thus $w_t = f_{i,t} = C$) has a unique steady state and is saddle-path stable. Label steady-state actions $\bar{A}$, the steady-state stock $\bar{S}$, and steady-state payoffs $\bar{\pi}$, so that $\bar{\pi} \triangleq \pi(C, \bar{A}, \bar{S}; K)$. I write $\bar{A}(K, C)$, so that $\mathrm{d}\bar{A}/\mathrm{d}C = \partial\bar{A}/\partial C + [\partial\bar{A}/\partial K][\mathrm{d}K/\mathrm{d}C]$. I assume henceforth that the stock begins at its deterministic steady state value (i.e., that $S_0 = \bar{S}$).

I first define the true effect of climate. I then analyze how past and future weather affect agents' choices. I finally consider an empirical researcher's ability to estimate the true effect of climate from variation in payoffs induced by past and future weather.

4.1 The True Effect of Climate on Payoffs

Following the empirical literature, we are interested in the long-run effects of altered climate on average payoffs. Appendix F.7 shows that

$$\begin{aligned}\lim_{t\to\infty} \frac{\mathrm{d}E_0[\pi_t]}{\mathrm{d}C} =& \bar{\pi}_w + \bar{\pi}_A \frac{\mathrm{d}\bar{A}}{\mathrm{d}C} + \bar{\pi}_S \frac{\mathrm{d}\bar{S}}{\mathrm{d}C} + \overbrace{\bar{\pi}_K \frac{\mathrm{d}K}{\mathrm{d}C}}^{=0} \\ =& \underbrace{\bar{\pi}_w}_{\text{direct effects}} + \underbrace{\left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g}\right] \frac{\mathrm{d}\bar{A}(K,C)}{\mathrm{d}C}}_{\text{adaptation effects}} .\end{aligned} \tag{10}$$

The direct effects of alterations to long-lived infrastructure K again vanish because agents optimize this infrastructure around long-run payoffs. However, adaptation choices A_t can now can have first-order consequences for average payoffs, both directly and through their effects on the stock. Their effects on the stock become increasingly important as the stock becomes more persistent (i.e., as g approaches 1).

Why do adaptation responses now have first-order effects on payoffs? In Section 3, changing these actions had no effect because the first-order condition ensured that $\pi_A = 0$. However, in a dynamic environment, agents set $V_A = 0$, not $\pi_A = 0$ (see equation (1)). Optimal actions satisfy the Euler equation, derived in Appendix F.5:

$$-\pi_A(w_t, A_t, S_t; K) = \beta h'(A_t)\, E_t\left[\pi_S(w_{t+1}, A_{t+1}, S_{t+1}; K) + g \frac{-\pi_A(w_{t+1}, A_{t+1}, S_{t+1}; K)}{h'(A_{t+1})}\right]. \tag{11}$$

Agents equate the marginal effect of actions on contemporary payoffs (the left-hand side) to the marginal effect of actions on expected future payoffs (the right-hand side), which includes the effect π_S of altering the stock and the effect of adjusting subsequent actions to return to the original stock trajectory. An agent may, for instance, choose an action whose marginal effect on immediate payoffs is negative if that action increases expected future payoffs. We recover the static efficiency condition that $\pi_A = 0$ only as agents become myopic ($\beta \to 0$) or as the stock becomes independent of actions ($h' \to 0$).[24]

Around the deterministic steady state, equation (11) implies (see (A-6)):

$$h'(\bar{A})\bar{\pi}_S = -\frac{1-\beta g}{\beta}\bar{\pi}_A.$$

Substitute into (10) to obtain:

$$\lim_{t\to\infty}\frac{\mathrm{d}E_0[\pi_t]}{\mathrm{d}C} = \bar{\pi}_w - \frac{1}{1-g}\frac{1-\beta}{\beta}\bar{\pi}_A\frac{\mathrm{d}\bar{A}(K,C)}{\mathrm{d}C}. \tag{12}$$

If the steady-state adaptation response to climate change increases per-period payoffs for a given stock (i.e., if $\bar{\pi}_A\frac{\mathrm{d}\bar{A}(K,C)}{\mathrm{d}C} > 0$), then it reduces steady-state payoffs. For instance, if the stock has no persistence ($g = 0$), the effect on the stock in (10) becomes $h'(\bar{A})\bar{\pi}_S$, which from the Euler equation (11) is equal to $-\bar{\pi}_A/\beta$. This term is larger than the effect $\bar{\pi}_A$ on contemporary payoffs in (10): a dynamically optimizing agent trades off short-run increases in per-period payoffs against long-run costs imposed through the stock, and discounting means those costs must be larger in current value terms. As agents become perfectly patient ($\beta \to 1$), the long-run costs are exactly offset by the short-run benefits, but as agents become myopic ($\beta \to 0$), those short-run benefits are obtained by imposing especially large costs on the future through a depleted stock. When we add the short-run benefits and long-run costs together in a steady-state calculation, the long-run costs dominate.

An empirical researcher will therefore need to estimate how climate affects actions around the deterministic steady state $\bar{A}$ if they are to recover the effect of climate on average payoffs. Appendix F.8 shows that

$$\frac{\mathrm{d}\bar{A}(K,C)}{\mathrm{d}C} \propto \underbrace{\overbrace{\bar{\pi}_{wA}}^{\text{ex-post adaptation}} + \overbrace{\beta\left[h'(\bar{A})\bar{\pi}_{wS} - g\,\bar{\pi}_{wA}\right]}^{\text{ex-ante adaptation}}}_{\propto\,\partial\bar{A}(K,C)/\partial C} + \overbrace{\underbrace{\left[(1-\beta g)\bar{\pi}_{AK} + \beta h'(\bar{A})\bar{\pi}_{SK}\right]}_{\propto\,\partial\bar{A}(K,C)/\partial K}\frac{\mathrm{d}K}{\mathrm{d}C}}^{\text{interactions with long-lived infrastructure}}. \tag{13}$$

There are three terms. The first captures what the literature has called reactive or ***ex-post adaptation*** to realized changes in weather (Fankhauser et al., 1999;

[24] But the presence of the stock makes optimized actions differ from the static setting of Section 3 even as $\beta \to 0$ or $h' \to 0$. We recover the same optimized actions as in the static setting only if, in addition, $\pi_{AS} = 0$, so that π_A is independent of the stock.

Mendelsohn, 2000). Here the reaction is to observed changes in weather induced by climate change. It depends on how weather shifts the marginal benefit of short-run actions, controlled by π_{wA}. For instance, farmers may water crops during a heat wave and firms may adjust production in response to price signals generated by weather events.

The second term captures what the literature has called anticipatory or ***ex-ante adaptation*** (Fankhauser et al., 1999; Mendelsohn, 2000). Here the anticipation reflects understanding of how climate change shifts all future weather. It depends on how weather shifts the marginal benefit of the stock, controlled by π_{wS}. For instance, farmers may conserve groundwater today in order to reduce the costs of irrigating in coming hot weather. Ex-ante adaptation also reflects agents anticipating that future actions will alter the stock in still-later periods. They therefore begin investing now to reduce distortions in the later stock. For instance, farmers may cut back on groundwater use today to make sure there is still enough groundwater left after the hot weather passes. Unsurprisingly, myopic agents ($\beta = 0$) do not undertake ex-ante adaptation.

The remaining terms depend on how long-lived infrastructure K responds to the change in climate. Changes in this infrastructure do not directly affect steady-state payoffs when optimized (i.e., $\bar{\pi}_K = 0$), but they do indirectly affect steady-state payoffs when the marginal benefit of either short-run actions ($\bar{\pi}_{AK}$) or the stock ($\bar{\pi}_{SK}$) depends on the choice of long-lived infrastructure. For instance, building irrigation canals might change the marginal cost of watering crops during a heat wave or the marginal benefit of having more groundwater.

4.2 How Weather Affects Decisions

An empirical researcher therefore needs to recover effects of climate on actions. I next build intuition for how weather determines actions in this environment before assessing whether the effects of climate change can be recovered from the effects of weather.

Figure 1 illustrates the determinants of time t actions. Formally, time t optimal actions are (Appendix F.5)

$$A_t = \bar{A} + \overbrace{\frac{\bar{\pi}_{wA}}{h'(\bar{A})\,\bar{\chi}}(w_t - C)}^{\text{effects of present weather}} + \overbrace{\bar{Z}(S_t - \bar{S})}^{\text{effects of past weather}} + \overbrace{\frac{\beta\Gamma}{h'(\bar{A})\,\bar{\chi}}\left[(f_{1,t} - C) + \frac{\beta\Psi}{h'(\bar{A})\bar{\chi}}(f_{2,t} - C)\right]}^{\text{effects of future weather}}, \tag{14}$$

where $h'(\bar{A})\,\bar{\chi} > 0$ and $\lim_{g\to 0} \bar{Z} \propto \bar{\pi}_{AS}$. The $\bar{\chi}$ and $\bar{Z}$ are functions of derivatives of $\bar{\pi}$. They result from a backward recursion that captures forward-looking optimization. I discuss Ψ below.

Present weather affects present actions through an ex-post adaptation channel. This channel is controlled by $\bar{\pi}_{wA}$, with actions aiming to mitigate the immediate harm

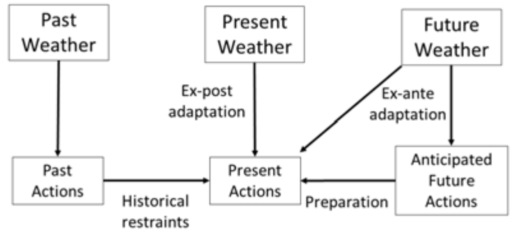

Figure 1: The determinants of present actions.

or amplify the immediate benefits of weather outcomes. This term is proportional to the ex-post adaptation channel in equation (13).

Past weather and forecasts affect present actions by altering the past actions that determine the present stock. The history of weather thereby restrains present actions. For g small, $\bar{Z}$ is proportional to $\bar{\pi}_{AS}$. When $\bar{\pi}_{AS} > 0$, past actions that increased the stock justify more present action, but when $\bar{\pi}_{AS} < 0$, past actions that increased the stock favor less present action.[25]

Future weather affects present actions through forecasts of that weather. The coefficients on forecasts in (14) are each proportional to $\beta\Gamma$, where

$$\Gamma \triangleq \overbrace{\left[h'(\bar{A})\bar{\pi}_{wS} - g\bar{\pi}_{wA}\right]}^{\text{ex-ante adaptation from (13)}} + \overbrace{\Psi\frac{\bar{\pi}_{wA}}{h'(\bar{A})\,\bar{\chi}}}^{\text{preparatory actions}} \tag{15}$$

and

$$\begin{aligned}\Psi \triangleq & h'(\bar{A})\bar{\pi}_{AS} + g\overbrace{\left(-\bar{\pi}_{AA} + \frac{h''(\bar{A})}{h'(\bar{A})}\bar{\pi}_A\right)}^{>0 \text{ by (3)}} \\ \propto & \left.\frac{\mathrm{d}A_t}{\mathrm{d}A_{t+1}}\right|_{w_t, f_{1,t}, f_{2,t}=C}.\end{aligned} \tag{16}$$

Unsurprisingly, the coefficients on forecasts in (14) go to zero as agents become myopic. For forward-looking agents, three terms in equation (15) control how actions depend on forecasts of future weather. First, when $\bar{\pi}_{wS} \neq 0$, agents choose today's actions in order to directly mitigate the consequences (or enhance the benefits) of expected future weather. This is the most direct form of ex-ante adaptation. Second, expecting higher weather outcomes in the future changes how agents trade-off time t and $t+1$ actions when trying to reach the desired time $t+2$ stock. If, for instance, a higher forecast makes future actions more valuable ($\bar{\pi}_{wA} > 0$), then agents

[25] Appendix F.5 shows that $\bar{Z} \to 0$ as either $h'(\bar{A})$ goes to zero or as $\bar{\pi}_{AS}$ and g jointly go to zero. As $h'(\bar{A}) \to 0$, past actions do not affect the stock around the steady state. As $\bar{\pi}_{AS} \to 0$, changes in the stock do not directly affect the marginal benefit of current actions, and as $g \to 0$, the time t stock does not affect the desired time $t+1$ stock or the time t actions taken to reach it.

cut back on current actions. This effect vanishes as $g \to 0$ because the time $t+2$ stock then depends only on time $t+1$ actions. This is an indirect form of ex-ante adaptation. These first two terms are proportional to the ex-ante adaptation channel in equation (13).

Third, agents anticipate how today's choices impose historical restraints on future choices and so undertake preparatory actions that can enable beneficial future actions. $\bar{\pi}_{wA}/[h'(\bar{A})\bar{\chi}]$ captures how a higher forecast shifts desired future actions. The term labeled Ψ captures how today's actions change with expectations of future actions. Equation (16) shows that Ψ depends on two terms. The first term within Ψ reflects intertemporal substitutability or complementarity among actions. When actions are intertemporal complements ($h'(\bar{A})\,\bar{\pi}_{AS} > 0$), a forecast that increases desired future actions leads agents to choose high actions today as a means of reducing future adjustment costs, but when actions are intertemporal substitutes ($h'(\bar{A})\,\bar{\pi}_{AS} < 0$), a forecast that increases desired future actions leads agents to choose low actions today as a means of conserving resources for the future. The second term within Ψ reflects how changes in desired future actions affect the tradeoff between time t and $t+1$ investments in reaching the desired time $t+2$ stock. This effect vanishes as $g \to 0$. The preparatory action term in equation (15) was absent from the effects of climate change derived in equation (13), a distinction that will be important for subsequent analysis.

4.3 Recovering the Effect of Climate from Weather Regressions

Now consider the possibility of estimating long-run climate impacts from variation in weather. By affecting people's lived experience of weather, a change in climate affects actions reactively chosen to deal with present weather. It also affects the past weather experienced by agents once they have been living in the counterfactual climate. This channel will make it important to estimate the effects of past weather. Finally, a change in climate affects agents' expectations of future weather, manifested as systematically higher forecasts. This channel will make it important to estimate the effects of forecasts.

I here assume that the empirical researcher can observe payoffs (e.g., profits) and weather variables.[26] I do not assume that the empirical researcher observes all of the actions that agents take or the level of the stock. Consider the following distributed lag regression with fixed effects:

$$\Delta\pi_{jt} = \alpha_j + \sum_{i=0}^{I} \Lambda_i \Delta w_{j(t-i)} + \sum_{i=0}^{I} \lambda_i \Delta f_{j1,(t-i)} + \sum_{i=0}^{I} \gamma_i \Delta f_{j2,(t-i)} + \eta_{jt}, \tag{17}$$

[26] The analysis straightforwardly extends to the case where the empirical researcher instead observes actions, not payoffs, and seeks the effect of climate on actions.

where units are again labeled j, where α_j is a fixed effect for agent j, where $I \geq 0$ controls the number of lags, and where η_{jt} is an error term. As before, a hat denotes estimated coefficients.

The proposition describes the effect of summing the estimated coefficients on an arbitrarily large number of lags:[27]

Proposition 2 (Summing Many Lags)**.**

$$\lim_{I\to\infty} \operatorname{plim} \sum_{i=0}^{I-3} \left[\hat{\Lambda}_i + \hat{\lambda}_i\right] = \bar{\pi}_w + \omega \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g}\right] \left(\frac{\partial \bar{A}(K,C)}{\partial C} + \Omega\right), \tag{18}$$

where

$$\Omega \propto \beta\Psi \frac{\bar{\pi}_{wA}}{h'(\bar{A})\,\bar{\chi}}$$

and Ψ *is defined in equation* (16)*. If* $\beta\Psi > 0$*, then* $\omega < 1$*. If* $\beta\Psi = 0$*, then* $\omega = 1$*. If* $\beta\Psi < 0$*, then* $\omega > 1$*.*

Proof. See Appendix F.9. Sketch: Expands π_t around the deterministic steady state, calculates each coefficient as the covariance of the associated variable with π_t divided by the variance of that variable, defines Ω and ω to reconcile the summed lags with the effect of climate change, and then analyzes these terms. □

The good news is that the right-hand side of (18) comes somewhat close to the true effect of climate derived in equation (10). In particular, it successfully captures the direct effect of weather and also captures effects proportional to ex-post and ex-ante adaptation.

However, three wedges make the true effect in equation (10) differ from the estimated effect in (18). First, the change in steady state actions in (18) holds K fixed, but equation (13) showed that $\mathrm{d}\bar{A}/\,\mathrm{d}C$ depends on changes in K when either $\bar{\pi}_{AK} \neq 0$ or $\bar{\pi}_{SK} \neq 0$. The problem is that long-lived infrastructure does not vary with weather shocks, so fluctuations in payoffs do not identify the consequences of adapting K to an altered climate. Even though these long-run adaptations do not have first-order consequences for payoffs when chosen optimally (i.e., $\bar{\pi}_K = 0$), equation (13) shows that these long-run adaptations can affect short-run actions that do have first-order consequences for payoffs. This wedge vanishes if long-lived infrastructure is in fact fixed over the timescale of climate change (i.e., if $\mathrm{d}K/\,\mathrm{d}C = 0$ in equation (13)) or if it does not directly interact with other decisions (i.e., if $\bar{\pi}_{AK} = \bar{\pi}_{SK} = 0$ in equation (13)).

[27] The requirement that we estimate at least I lags even though we use only $I - 3$ lags in the calculation avoids ancillary complications from omitted variables bias at the longest lags used in the calculation. If regression (17) were in levels, we would obtain the same result with one less lag required on the right-hand side.

The remaining two wedges arise from the durability of shorter-run decisions. Ω is a bias in estimated ex-ante adaptation. It is proportional to the preparatory actions defined in equation (15). Ex-ante adaptation is identified from transient shocks to forecasts. Preparatory actions reflect that an idiosyncratically high forecast implies idiosyncratically high future weather, for which current actions are not the most suited. An increase in the climate index C also increases forecasts but does so systematically rather than idiosyncratically: because increasing C also increases current and past weather, preparations for a change in weather are not relevant to the long-run effects of climate. Forecasts are critical to identifying ex-ante adaptation, but agents do not respond to higher-than-average forecasts in exactly the same way as they respond to forecasts that reflect higher average weather.[28]

The final wedge is ω. This term reflects the difference between the historical restraints on current actions imposed by transient weather shocks and those imposed by a change in climate that affects all past weather realizations and all past forecasts. By the proposition and (16), we have $\omega < 1$ when actions are intertemporal complements. In this case, historical restraints prevent an agent from adjusting too much to any particular transient weather shock, but when that shock has been repeated many times in the past (as eventually happens following a change in climate), the many small adjustments eventually add up to a much greater adjustment. Responses to transient shocks therefore overstate historical restraints, which is why $\omega < 1$. Consistent with conjectures in Deschênes and Greenstone (2007), observable short-run adaptation is less than long-run adaptation.

In contrast, we can have $\omega > 1$ when actions are intertemporal substitutes. In this case, an agent can experience more severe historical restraints following a change in climate than following a transient weather shock. For instance, if actions depend on scarce resources, agents may respond strongly to a transient weather shock but be unable to maintain this response for a long period of time. Their response to a change in climate may thus be relatively muted. Responses to transient shocks can understate historical restraints, which is why $\omega > 1$. Consistent with conjectures in Fisher et al. (2012) and Blanc and Schlenker (2017), observable short-run adaptation is greater than long-run adaptation.[29]

Proposition 2 described the results of estimating a model with infinite lags and summing the coefficients. The following corollary describes feasible regressions, with

[28]One could eliminate Ω by not using the forecast coefficients $\hat{\lambda}_i$, instead relying on $\lim_{I\to\infty}\sum_{i=0}^{I-2}\hat{\Lambda}_i$. However, this calculation would introduce a new bias, as it would miss all ex-ante adaptation terms in equation (13). One might also consider including additional forecast horizons in the summation. Summing the first and second horizons multiplies the ex-ante adaptation component and Ω by $1 + \beta\Psi/[h'(\bar{A})\bar{\chi}]$, introducing a new bias. If we had infinite forecast horizons, summing them would multiply the ex-ante adaptation component and Ω by $1/\{1 - \beta\Psi/[h'(\bar{A})\bar{\chi}]\}$, again introducing a new bias. Neither formulation clearly improves on (18).

[29]If $g > 0$, we can have $\omega < 1$ even when actions are intertemporal substitutes (via the second term in (16)). The reason is that an agent living in an altered climate may choose actions to loosen historical restraints over time.

fewer lags:[30]

Corollary 3 (Summing Finite Lags)**.** *Let $I' \geq 1$ and $I \geq I' + 3$. Then:*

$$\operatorname{plim} \sum_{i=0}^{I'} \left[\hat{\Lambda}_i + \hat{\lambda}_i\right] = \bar{\pi}_w + \omega_{I'} \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g}\right] \left(\frac{\partial \bar{A}(K,C)}{\partial C} + \Omega\right),$$

where Ω is as in Proposition 2. If $\Psi = 0$, then $\omega_{I'} = \omega = 1$. If $\Psi > 0$, then $\omega_{I'} \in (0, \omega)$ with $\omega < 1$ and $\omega_{I'}$ increasing in I'. If $\Psi < 0$, then $\omega_{I'} > \omega > 1$ for I' odd.

Proof. See Appendix F.10. □

The number of summed lags affects only ω. When $\omega < 1$, responses to weather shocks underestimate responses to long-run changes in climate. Corollary 3 shows that this underestimation is more severe when based on a shorter history of weather shocks. Matters are more complicated when $\omega > 1$. In this case, the bias $\omega_{I'}$ fluctuates as we increase I', clearly introducing more bias than ω when I' is odd.[31]

The net bias introduced by the wedges Ω and ω cannot be signed in general. However, both wedges do vanish in some intuitive special cases, leaving only the wedge induced by K being fixed. The first special case is when agents are myopic:

Corollary 4 (Myopic Agents)**.**

$$\lim_{\beta \to 0} \lim_{I \to \infty} \operatorname{plim} \sum_{i=0}^{I-2} \hat{\Lambda}_i = \bar{\pi}_w + \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g}\right] \frac{\partial \bar{A}(K,C)}{\partial C}.$$

Proof. See Appendix F.11. □

First, in a special case with myopic agents who do not undertake ex-ante adaptation ($\beta = 0$), the wedge introduced by preparatory actions vanishes because myopic agents are not concerned about future actions. The sign of the bias then depends only on the wedge ω induced by historical restraints, as even myopic agents respond to their own past decisions (see also Keane and Wolpin, 2002). And this wedge vanishes as we sum an infinite number of lags: myopic agents respond to a long sequence of transient weather shocks in exactly the same way as they respond to living in a world with an

[30] The requirement that we estimate at least I lags even if we use only I' lags avoids complications from omitted variables bias at the longest lags used in the calculation. See footnote 27.

[31] Appendix A analyzes "long difference" estimators, which average over δ timesteps and estimate a conventional weather regression on the transformed data (e.g., Dell et al., 2012; Burke and Emerick, 2016). Long difference estimators are motivated by the possibility that climate change has manifested itself over long timesteps, but Appendix A shows that long difference estimators are identified by sequences of transient weather shocks even when the climate has been constant. At best, these estimators conflate the two sources of variation, and at worst they are identified off nothing but the transient weather shocks. In the latter case, Appendix A shows that long difference estimators are inferior to simply estimating regression (17) with $I \geq \delta + 3$ lags.

altered climate.[32] Therefore we recover the effect of climate by estimating a regression with arbitrarily many lags if agents are myopic and long-lived infrastructure either is fixed or does not interact with shorter-run adaptation decisions.[33]

The next special case makes actions independent of each other over time:

Corollary 5 (Independent Actions)**.** *Let $I' \geq 1$ and $I \geq I'+3$. If $\pi_{AS}(\cdot, A_t, S_t; K) = 0$, then:*

$$\lim_{g \to 0} \operatorname{plim} \sum_{i=0}^{I'} \left[\hat{\Lambda}_i + \hat{\lambda}_i\right] = \bar{\pi}_w + \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g}\right] \frac{\partial \bar{A}(K,C)}{\partial C}.$$

Proof. See Appendix F.12. □

Each period's decisions are independent of other periods' decisions in a special case without interactions between different periods' actions ($\bar{\pi}_{AS}, g = 0$). In equation (14), we lose the effects of past weather (see footnote 25). Estimating effects of realized weather suffices to recover the direct effects of climate as well as the effects of ex-post adaptation, and estimating effects of forecasts suffices to recover the effects of ex-ante adaptation. In fact, in this special case we do not even need to estimate all of the lags. When actions are chosen independently over time, the coefficients on lags longer than the first are all zero. These can be dropped from the regression without causing bias. But it is still important to include the first lag of both weather and forecasts. This lag picks up effects of time $t-1$ weather and forecasts on time t payoffs, via the effects of time $t-1$ actions on the time t stock. In equation (10), the contemporary effects identify $\bar{\pi}_w + \bar{\pi}_A \frac{\partial \bar{A}(K,C)}{\partial C}$ and the lagged effects identify $h'(\bar{A})\bar{\pi}_S \frac{\partial \bar{A}(K,C)}{\partial C}$. Therefore we recover the effect of climate by estimating a regression with at least one lag of weather and forecasts if $\bar{\pi}_{AS}, g = 0$ and long-lived infrastructure either is fixed or does not interact with shorter-run adaptation decisions.

5 An Indirect Least Squares Estimator of Climate Impacts

We have thus far seen that we can exactly recover the effects of climate change from simple weather regressions only under restrictive assumptions: if agents are not affected by resource or capital stocks, if agents are myopic and long-lived infrastructure

[32] The bias introduced by $\omega_{I'}$ in Corollary 3 does not vanish as $\beta \to 0$: even myopic agents respond to the weather they lived through and experience the historical restraints imposed by their responses. Only by estimating infinite lags of weather can we replicate the long-run effect of living in an altered climate.

[33] When agents are myopic, we do not need to estimate responses to forecasts (and should obtain $\operatorname{plim} \hat{\lambda}_i = 0$ and $\operatorname{plim} \hat{\gamma}_i = 0$ if we do). On the other hand, when agents are not myopic but forecasts do not exist, we lose the bias induced by preparatory actions but we also fail to estimate the ex-ante adaptation terms from (13), as even agents who lack short-run forecasts may respond proactively to their knowledge of climate change.

either is fixed or does not interact with shorter-run adaptation decisions, or if agents make decisions independently over time and long-lived infrastructure either is fixed or does not interact with shorter-run adaptation decisions. But although we have described the biases that arise when these conditions do not hold, we have not been able to sign any parts of those biases.

I now show how an indirect least squares estimator can disentangle direct effects from each type of adaptation and bound climate impacts, modulo only effects of altered long-lived infrastructure. I derive this new estimator by solving for the probability limit of each estimated coefficient in terms of model primitives and then rearranging those equations to solve for the combinations of model primitives needed to calculate climate change impacts. Importantly, this new approach maintains precisely the same credible identification as the reduced-form specifications in Section 4.

I derive the estimator first in the presence of forecasts and subsequently in the absence of forecasts.

5.1 In the Presence of Forecasts

Begin by considering the case we have analyzed so far, in which agents see specialized forecasts of coming weather and the empirical researcher has data on those forecasts. For example, in Shrader (2020), the empirical researcher observes forecasts of the El Niño seasonal climate index. The following proposition presents the indirect least squares estimator for climate impacts:

Proposition 6 (Indirect Least Squares, With Forecasts)**.** *Consider estimating regression* (17)*, with* $\operatorname{plim} \hat{\lambda}_0, \hat{\lambda}_1, \hat{\Lambda}_1 \neq 0$*. Then, for* $I > 3$*,*

$$\begin{aligned}
&\operatorname{plim}\left(\overbrace{\hat{\Lambda}_0 - \hat{\Lambda}_1\frac{\hat{\lambda}_0}{\hat{\lambda}_1}}^{\textit{direct effects}} \overbrace{- \frac{1-\beta}{\beta}\hat{\Lambda}_1\frac{\hat{\lambda}_0}{\hat{\lambda}_1}}^{\textit{ex-post adaptation}} \overbrace{\overbrace{- \frac{1-\beta}{\beta}\hat{\lambda}_0}^{\textit{estimated}} \; \overbrace{+ \frac{1-\beta}{\beta}\frac{\hat{\gamma}_0}{\hat{\lambda}_0}\hat{\Lambda}_1\frac{\hat{\lambda}_0}{\hat{\lambda}_1}}^{\textit{prep. action adjustment}}}^{\textit{ex-ante adaptation}} \right) \\
=&\bar{\pi}_w + \tilde{\omega}\left[\bar{\pi}_A + \bar{\pi}_S\frac{h'(\bar{A})}{1-g}\right]\frac{\partial \bar{A}(K,C)}{\partial C},
\end{aligned} \tag{19}$$

$$\operatorname{plim}\frac{\hat{\lambda}_0}{\hat{\lambda}_1} < 0,$$

and

$$\Psi \propto \operatorname{plim}\frac{\hat{\Lambda}_2}{\hat{\Lambda}_1}.$$

If $\Psi > 0$*, then* $\tilde{\omega} < 1$*. If* $\Psi = 0$*, then* $\tilde{\omega} = 1$*. If* $\Psi < 0$*, then* $\tilde{\omega} > 1$*.*

Proof. See Appendix F.13. Sketch: Follows proof of Proposition 2 to derive the probability limits of estimated coefficients and then solves the system for the desired terms. □

The indirect least squares estimator provides an estimate of the effect of climate change and also provides estimates of the channels that determine that effect.[34] Comparing the right-hand side of (19) to (10), we see the usual bias due to K being fixed and an additional bias when $\tilde{\omega} \neq 1$ that is signed from estimated coefficients.

Consider how the new estimator identifies direct effects separately from ex-post adaptation. The estimated coefficient $\hat{\Lambda}_0$ on contemporary weather identifies the sum of direct weather effects and the immediate payoffs from ex-post adaptation to that weather. However, as discussed around (12), the steady-state effects of ex-post adaptation to climate change must account for its dynamic consequences, not just its immediate payoffs. If, for instance, ex-post adaptation provides short-run benefits at the expense of long-run costs, the coefficient $\hat{\Lambda}_1$ on lagged weather identifies those long-run costs. Past forecasts affect current payoffs only through the same types of dynamic tradeoffs that determine $\hat{\Lambda}_1$, so the ratio $\hat{\lambda}_0/\hat{\lambda}_1$ uses variation in forecast timing to identify the ratio of short-run to long-run effects that also determines optimal ex-post adaptation. Multiplying $\hat{\Lambda}_1$ by $\hat{\lambda}_0/\hat{\lambda}_1$ converts, for instance, dynamic costs of optimally chosen actions to immediate benefits. Subtracting $(\hat{\lambda}_0/\hat{\lambda}_1)\hat{\Lambda}_1$ from $\hat{\Lambda}_0$ then eliminates the immediate payoffs of ex-post adaptation from $\hat{\Lambda}_0$, leaving only the direct effects of weather. And by converting the dynamic effects of long-run adaptation into a steady-state effect (via $(1-\beta)/\beta$), we isolate the long-run effect of ex-post adaptation to climate change.

Now consider how the new estimator adjusts the ex-ante adaptation channel to remove the effects of preparatory actions. The estimated coefficient $\hat{\lambda}_0$ on contemporary forecasts identifies a term related to ex-ante adaptation. But that term also includes preparatory actions irrelevant to changes in climate (see discussions of Ψ in (15) and Ω in (18)). Fortunately, longer-horizon forecasts identify preparatory actions, as there is no other channel through which longer-horizon forecasts can affect current payoffs in this environment. Formally, the coefficient $\hat{\gamma}_0$ on longer-horizon forecasts contains Ψ from equation (14), Ψ is the term that drives preparatory actions in (15), and the adjustment to $\hat{\gamma}_0$ accounts for effects such as $\bar{\pi}_{wA}$ in (15).[35] The indirect least squares estimator adjusts the ex-ante adaptation channel to remove the preparatory action bias.

Finally, the ratio of coefficients on lagged weather identifies how actions are linked over time. As mentioned above, past weather matters for current weather via ex-post

[34] The indirect least squares estimator defined in Proposition 6 is a consistent estimator of the indicated combinations of theoretical primitives, but because it is a nonlinear function of the estimated ordinary least squares coefficients, it is not unbiased. A similar comment will apply to subsequent results.

[35] That adjustment is reminiscent of the ex-post adaptation term because $\bar{\pi}_{wA}$ defines ex-post adaptation (see (13)).

adaptation, which in turn matters for current payoffs via actions' dynamic consequences. The sign of Ψ controls the bias from $\tilde{\omega}$, as it did for ω in (18), and we can infer the sign of Ψ from $\hat{\Lambda}_2/\hat{\Lambda}_1$.[36] We thus learn from the ratio of lagged weather coefficients whether $\tilde{\omega}$ dampens or inflates the adaptation channels.

There are two cases. First, if the ratio of coefficients on lagged weather is positive, then $\Psi > 0$ and $\tilde{\omega} < 1$. Adaptation to climate is greater than implied by responses to weather, as when adjustment costs constrain short-run responses more than long-run responses (i.e., actions are intertemporal complements).[37] In that case, the top line of (19) gives a lower (upper) bound on the true effect of climate if the adaptation terms are positive (negative). Because adaptation could be arbitrarily large, we have only a one-sided bound. If, instead, the ratio of coefficients on lagged weather is negative, then $\Psi < 0$ and $\tilde{\omega} > 1$. Adaptation to climate is less than implied by responses to weather, as when resource constraints bind in the long run but not in the short run (i.e., actions are intertemporal substitutes). In that case, the top line of (19) and the estimated direct effects bound the effect of climate from either side.

Either way, we have bounded the effect of climate as long as long-lived infrastructure either is fixed or does not interact with shorter-run adaptation decisions. And remarkably, we have done so while maintaining reduced-form identification, without needing to observe either the stock or actions, and without needing to estimate the payoff function or the stock accumulation equations. In contrast, the simpler ordinary least squares calculations in Sections 3 and 4.3 do not generally bound the effects of climate and can exactly recover the effects of climate only when the decision-making environment is sufficiently simple: Proposition 1 required $\pi_S = 0$, and Corollaries 4 and 5 required either that agents are myopic ($\beta = 0$) or that actions are independent over time ($g, \bar{\pi}_{AS} = 0$), in addition to long-lived infrastructure either being fixed or not interacting with shorter-run adaptation decisions. Of course, the present section's calculations also exactly recover the effects of climate when these special conditions are met. The indirect least squares approach therefore directly weakens the prerequisites to recover climate from weather, without sacrificing anything in terms of identification.

5.2 In the Absence of Forecasts

What if agents do not have specialized forecasts of future weather, so they predict future weather solely from their climate zone? For instance, farmers may lack quality forecasts of temperature and rainfall months or more ahead. In this case, effects of future weather drop out of (14). However, the lack of forecasts does not absolve empirical researchers from needing to estimate the ex-ante adaptation to climate

[36] Deschênes and Greenstone (2012) noted the importance of estimating lags when thinking about actions such as storage.

[37] Although note, from (16), that if $g > 0$, then intertemporal substitutes can be consistent with positive Ψ.

change described in (13): whether or not agents have specialized forecasts of future weather, they can recognize when the climate has changed and can understand what a changed climate means for average weather.

Consider the following regression, which is similar to (17) but omits the nonexistent forecasts:

$$\Delta\pi_{jt} = \alpha_j + \sum_{i=0}^{I} \phi_i \Delta w_{j(t-i)} + \eta_{jt}. \tag{20}$$

The following proposition presents the indirect least squares estimator for climate impacts in this setting without forecasts:

Proposition 7 (Indirect Least Squares, Without Forecasts). *Consider estimating regression* (20)*, assuming that* $\operatorname{plim}\hat{\phi}_1 \neq 0$*. For* $I > 3$*,*

$$\operatorname{plim}\left(\overbrace{\hat{\phi}_0 + \frac{\hat{\phi}_1}{\frac{1}{\beta} - \frac{\hat{\phi}_2}{\hat{\phi}_1}}}^{\text{direct effects}} + \overbrace{\frac{1-\beta}{\beta}\frac{\hat{\phi}_1}{\frac{1}{\beta} - \frac{\hat{\phi}_2}{\hat{\phi}_1}}}^{\text{ex-post adaptation}}\right) = \bar{\pi}_w + \tilde{\omega}\left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g}\right] \lim_{\beta \to 0} \frac{\partial \bar{A}(K,C)}{\partial C},$$

$$\left|\operatorname{plim}\frac{\hat{\phi}_2}{\hat{\phi}_1}\right| < 1,$$

and

$$\Psi \propto \operatorname{plim}\frac{\hat{\phi}_2}{\hat{\phi}_1}.$$

If $\Psi > 0$*, then* $\tilde{\omega} < 1$*. If* $\Psi = 0$*, then* $\tilde{\omega} = 1$*. If* $\Psi < 0$*, then* $\tilde{\omega} > 1$*.*

Proof. See Appendix F.14. Sketch: Follows proof of Proposition 2 to derive estimated coefficients and then solves the system for the desired terms. □

The intuition for identification is much as given following Proposition 6. The intuition for the adjustment to obtain direct effects is also as described following Proposition 6. In the absence of observable forecasts, the adjustment for the dynamic consequences of ex-post adaptation depends on how the ratio of lagged weather impacts compares to the discount factor β. The biases introduced by the possibility of endogenous K and by $\tilde{\omega} \neq 1$ are also as in Proposition 6. However, we now have a new bias due to the inability to estimate ex-ante adaptation, which means that we recover effects on actions only as agents become myopic (hence the $\lim_{\beta\to 0}$ on the right-hand side). Nothing in the data can identify ex-ante adaptation without variation in agents' expectations of weather.

6 Application: Impacts of Climate Change on the U.S. Economy

I demonstrate the new indirect least squares approach by analyzing how temperature affects output and income in U.S counties. A competitive economy acts to maximize the present discounted value of the representative agent's utility payoffs over time. The theoretical analysis has agents treat climate change as exogenous, which is consistent with a constrained-efficient economy that achieves an efficient allocation around a given climate trajectory and treats the effect of that allocation on climate change as an unpriced externality. For the application, I assume the representative agent has a logarithmic utility function over either output per capita or income per capita.

6.1 Data and Methods

I obtain county-level output and income per capita for the continental U.S. states from the Bureau of Economic Analysis. Output data cover 2001–2019, and income per capita data cover 1969–2019. I derive output per capita by using population data available from the NIH Surveillance, Epidemiology, and End Results (SEER) Program. Weather data follow Schlenker and Roberts (2006, 2009).[38] I aggregate weather on a 4km-by-4km grid to county-average weather. Appendix B reports descriptive statistics. I calibrate β to a 15% annual discount rate in the base specification.

Quality forecasts are unlikely to be available to economic agents at the timescale of these data. The analysis of Section 5.2 is therefore the most relevant to this application. Adapting regression (20), I estimate:

$$\Delta\pi_{jt} = \sum_{k=1}^{4}\sum_{i=0}^{4}\phi_{irk}\Delta w_{j(t-i)k} + \xi_{s1}t + \xi_{s2}t^2 + \alpha_j + \nu_{rt} + \eta_{jt}, \tag{21}$$

where π_{jt} is either log output per capita or log income per capita (i.e., the representative agent's flow utility) in county j and year t and where county j is in Census region r. Following prior work (e.g., Newell et al., 2021), I difference variables to minimize the effects of trends. Weather variable k counts the number of days in year t on which county j's daily average temperature falls in bin k. Following Leduc and Wilson (2023), the first bin counts days below 20°F, the second bin counts days in 20–50°F, the third bin counts days in 70–80°F, and the fourth bin counts days above

[38] I first estimate the relationship between a stable set of weather station observations from the NOAA Global Historical Climatology Network and 4km-by-4km weather interpolations from the Parameter-Elevation Regressions on Independent Slopes Model (PRISM). I then use the estimated relationship to aggregate weather station observations to a 4km-by-4km grid. This method creates a record of daily data that accounts for the physical characteristics used by PRISM while also limiting the scope for nonclassical measurement error induced by endogenous entry of weather stations over time.

80°F. The omitted bin is days in 50–70°F. Following the focus of much prior literature, I focus on the estimate for the hottest bin in the main text. Appendix E reports results for other bins.

I allow the effects of weather to vary by Census region r in order to mitigate concerns about heterogeneous treatment effects in linear fixed effects estimators (see Carter et al., 2018). I implement this heterogeneity by estimating separate regressions by Census region. The quadratic time trends vary by state, indexed by s. The α_j are county fixed effects, and the ν_{rt} are region-year fixed effects. In the base specification, regressions are weighted by each county's log population in 2002, which downweights outlier counties that contain very few people (in particular, some oil-rich rural counties).

Regression (21) includes four lags of weather, following Proposition 7. In order to compare to conventional approaches, I also adapt regression (8) so as to obtain estimates without lags of weather:

$$\Delta\pi_{jt} = \sum_{k=1}^{4} \theta_{rk} \Delta w_{jtk} + \xi_{s1} t + \xi_{s2} t^2 + \alpha_j + \nu_{rt} + \eta_{jt}. \tag{22}$$

I use the estimates from this regression to project climate impacts in the fashion that has been standard in the literature and that follows Proposition 1. Weighting and other specification choices follow regression (21).

6.2 Results

Table 1 reports the estimated coefficients on days above 80°F for regressions (21) and (22). Standard errors, in parentheses, are clustered by county, which accounts for unobserved correlation within a county over time.[39] In regression (22), more days of extreme heat significantly reduce both output per capita and income per capita in the Midwest and South and have insignificant effects in the West. Including lags makes contemporary extreme heat appear more harmful in the Midwest and Northeast, less harmful in the South, and potentially beneficial in the West.

In Section 3, I showed that conventional calculations of climate impacts from regression (22) require $\pi_S = 0$ in order to be informative about climate change impacts. However, the results of regression (21) contradict $\pi_S = 0$. First, lagged temperatures appear to affect both output per capita and income per capita, as several lagged weather variables have coefficients that are significant at the 1% level. Past weather presumably matters by affecting stock variables that subsequently affect payoffs. Second, the bottom panel of Table 1 shows that a Wald test rejects the hypothesis that

[39] This level of clustering is consistent with the bootstrapped confidence intervals discussed below. Two-way clustering is not straightforward to bootstrap. There are too few states in a Census region to cluster by state. Appendix D reports results with state-year clustering and without clustering.

Table 1: Estimating effects of days above 80 degF on county output per capita.

	Output p.c.		Income p.c.	
	Regression (21)	Regression (22)	Regression (21)	Regression (22)
Midwest				
Contemporary	-0.0013***	-0.00080***	-0.0011***	-0.00085***
	(0.00018)	(0.00014)	(0.000075)	(0.000071)
Lag 1	-0.00069***		-0.00033***	
	(0.00021)		(0.000073)	
Lag 2	0.00010		0.00015**	
	(0.00023)		(0.000071)	
Northeast				
Contemporary	-0.00059**	-0.00027*	-0.00025***	-0.000088
	(0.00024)	(0.00014)	(0.000077)	(0.000054)
Lag 1	-0.00038		-0.00029***	
	(0.00029)		(0.00010)	
Lag 2	-0.00022		-0.00032**	
	(0.00036)		(0.00013)	
South				
Contemporary	-0.0000021	-0.00015**	-0.0000088	-0.00012***
	(0.000079)	(0.000061)	(0.000034)	(0.000029)
Lag 1	0.00036***		0.00019***	
	(0.000096)		(0.000032)	
Lag 2	-0.000051		0.000077**	
	(0.00011)		(0.000032)	
West				
Contemporary	0.00052**	0.00017	0.000068	-0.000032
	(0.00024)	(0.00019)	(0.00011)	(0.000095)
Lag 1	0.00034		0.00014	
	(0.00028)		(0.00011)	
Lag 2	-0.00018		0.00018*	
	(0.00024)		(0.000092)	
Wald test of the null that the first two lags are jointly zero				
p-value, Midwest	0.000023		1.1e-12	
p-value, Northeast	0.32		0.017	
p-value, South	0.00000061		7.6e-10	
p-value, West	0.16		0.15	

Standard errors in parentheses
Standard errors clustered by county.
Regressions weighted by each county's log population in 2002.
Years: 2002–2019 for output and 1970–2019 for income.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: Indirect least squares estimates for the effects of days above 80 degF, from Proposition 7 and regression (21).

	Direct Effects	Ex-Post Adaptation	Total Effects	$\hat{\phi}_{2r4}/\hat{\phi}_{1r4}$
		Output p.c.		
Midwest	-0.0019	-0.000080	-0.0019	-0.15
	(-0.0028,-0.0012)	(-0.00020,-0.000017)	(-0.0030,-0.0012)	(-1.47,0.40)
Northeast	-0.0013	-0.00010	-0.0014	0.58
	(-0.029,0.0017)	(-0.0048,0.00025)	(-0.033,0.0018)	(-9.09,3.40)
South	0.00028	0.000042	0.00032	-0.14
	(-0.000042,0.00074)	(0.000013,0.000100)	(-0.000032,0.00083)	(-1.18,0.40)
West	0.00072	0.000030	0.00075	-0.54
	(-0.000081,0.0018)	(-0.000048,0.00015)	(-0.000100,0.0019)	(-55.2,2.16)
		Income p.c.		
Midwest	-0.0013	-0.000031	-0.0014	-0.46
	(-0.0016,-0.0011)	(-0.000060,-0.000013)	(-0.0017,-0.0011)	(-1.23,0.029)
Northeast	-0.0077	-0.0011	-0.0089	1.11
	(-1.35,-0.0043)	(-0.20,-0.00060)	(-1.55,-0.0049)	(0.54,2.32)
South	0.00025	0.000039	0.00029	0.39
	(0.000068,0.00050)	(0.000020,0.000068)	(0.000091,0.00057)	(0.073,0.66)
West	-0.0017	-0.00026	-0.0020	1.23
	(-0.21,-0.00068)	(-0.031,-0.00012)	(-0.24,-0.00078)	(-2.17,29.8)

95% confidence intervals (in parentheses) bootstrapped from 1000 samples with resampling at the county level.
Regressions weighted by each county's log population in 2002.
Years: 2002–2019 for output and 1970–2019 for income.

the first two lags are jointly zero in the Midwest and the South for either dependent variable and in the Northeast for income per capita, and it nearly rejects at conventional significance thresholds in the West.

The decision-making environment appears to be dynamic, so conventional calculations do not have a clear interpretation in terms of climate change. But the indirect least squares (ILS) estimator is valid in the presence of dynamics. Table 2 uses the estimated $\hat{\phi}$ from regression (21) in the formulas derived in Proposition 7. The 95% confidence intervals (in parentheses) are bootstrapped, with resampling at the county level. The ILS estimator's use of estimated coefficients in denominators makes the confidence intervals asymmetric.

The first column of Table 2 shows that the estimated direct effects of extreme heat are more deleterious in the Midwest and Northeast than what one might have expected from the coefficients on contemporary weather reported in Table 1: an extra day with an average temperature above 80°F in average temperature has the direct effect of reducing output (income) per capita by 0.19% (0.13%) in the Midwest and by 0.13% (0.77%) in the Northeast, whereas the coefficients on contemporary temperature from regression (21) implied reductions of only 0.13% (0.11%) and 0.06% (0.03%), respectively (and the coefficients from regression (22) implied even smaller effects). As described following Proposition 6, the reason for the gap is that the reduced-form coefficients in Table 1 entwine direct effects of temperature with ex-post adaptation that responds to temperature. In both of these regions, coefficients on contemporary and lagged extreme heat ($\hat{\phi}_{0r4}$ and $\hat{\phi}_{1r4}$) are negative. The negative coefficient on lagged extreme heat means that ex-post adaptation imposes dynamic costs, which an optimizing agent must have traded off against short-run benefits. Subtracting these short-run benefits from the coefficient on contemporary extreme heat isolates the direct effects of weather.

Estimating dynamic costs of ex-post adaptation also implies that ex-post adaptation further reduces both output and income per capita by around 0.01% per day of extreme heat in steady state (second column of Table 2). Combining the steady-state costs of direct effects and ex-post adaptation (third column), an extra day of extreme heat reduces output (income) per capita in the Midwest by 0.19% (0.14%), with the 95% confidence interval extending from a reduction of 0.28% (0.17%) to a reduction of 0.12% (0.11%). Either point estimate is around twice as large as the point estimate of a 0.08% (0.09%) reduction from regression (22). Moreover, neither confidence interval for total effects overlaps with the corresponding confidence interval from regression (22). The ILS estimator makes an extra day of extreme heat due to climate change appear significantly more harmful (in both practical and statistical senses) in the Midwest than suggested by the short-run effect of an extra day of extreme heat.

The ILS estimates for the Northeast are much less precise. The central estimate for output (income) per capita indicates a reduction of 0.14% (0.89%) per day of extreme heat, but the 95% confidence interval extends from a reduction of 3.3% (155%) to an increase of 0.18% (reduction of 0.49%). The imprecision in the Northeast reflects that the estimated effects of lagged weather are similar across the first and second lags in the Northeast (see Table 1), which makes the denominators in Proposition 7 small.

The central estimates in the South are for a benefit of extreme heat, but these estimates are small in magnitude (0.03% for both output per capita and income per capita). Whereas the estimated effects on output and income per capita cohere in most regions, they differ in the West. There, extreme heat is estimated to have a beneficial (albeit insignificant) effect on output per capita but a significantly harmful effect on income per capita, of comparable magnitude to impacts in the Midwest and Northeast.

The final column shows that estimates for $\hat{\phi}_{2r4}/\hat{\phi}_{1r4}$, whose sign determines the direction of bias when projecting ex-post adaptation (see Proposition 7). For output per capita, the short panel makes it hard to tell whether the second lag of extreme heat has a different sign from the first lag. But for income per capita, the estimated ratio is significantly positive in both the Northeast and South (and, in accord with the theory, the confidence intervals include values strictly less than 1). The positive sign suggests that adaptation will be greater in the long-run than in the short-run, as when actions are intertemporal complements (see equation (16)). In this case, steady-state ex-post adaptation will be greater than projected here, which would make climate change more beneficial in the South and more harmful in the other three regions.

Appendix C analyzes effects on output by industry rather than by region. It shows that the ILS and OLS estimators predict similar, small effects of extreme heat in most industries. There are two exceptions: the ILS estimator predicts steady-state losses in agriculture and retail that are much larger than—and significantly different from—those predicted by the OLS estimator. Estimating significant effects of extreme heat on agriculture accords with estimating significant effects of extreme heat in the Midwest. In each of these industries, the ratio $\hat{\phi}_{2r4}/\hat{\phi}_{1r4}$ is significantly positive. This sign again suggests that ex-post adaptation will be greater in the long-run than in the short-run and thus greater than predicted by the ILS estimator.

Appendix D reports robustness checks that vary handling of time trends, weighting, the number of lags included, clustering, and the discount rate. The main takeaways are fairly robust: extreme heat is harmful in the Midwest and (for income per capita) the West, ILS estimates in the Northeast are imprecise, estimated effects in the South are small, and the OLS estimator underestimates the steady-state harm from regular extreme heat. Including more lags makes estimated effects in the Midwest even larger. Estimated effects in the Northeast are sensitive to including time trends, to the sample of years, and to the choice of discount rate. Estimated effects on income per capita in the West are sensitive to the weighting scheme. Clustering by state-year instead of county makes estimated effects on output per capita in the Midwest and income per capita in the South no longer significant but does not affect the significance of harm to income per capita in the Midwest, Northeast, and West.

6.3 Climate Change Impacts

I project the effects of end-of-century climate change on U.S. output in Table 3. I use 30 models in the Coupled Model Intercomparison Project Phase 6 (CMIP6), downscaled by Thrasher et al. (2022). Following recommendations in Hausfather and Peters (2020), I use RCP 4.5 as a scenario with weak additional mitigation. I obtain the predicted change in days within each weather bin for each county for 2081–2100 relative to 1995–2014. I obtain projected changes in each Census region by weighting counties' changes by their log population in 2019. As reported in the notes to Table 3, the projected changes in annual days above 80°F range from an increase of roughly 2 days per year in the Northeast and West to an increase of 12 days per year in the Midwest and 21 days per year in the South.

Table 3 multiplies these projected changes in weather by the ILS and OLS estimates for the total effects of each weather variable in each region. The first column under each estimator shows projections for extreme heat. The point estimates suggest losses from additional extreme heat in the Midwest and Northeast, whether in terms of output per capita or in terms of income per capita. The ILS estimator projects losses that are significantly larger than projected by the OLS estimator, in both practical and statistical terms. In particular, the ILS estimator projects losses to income per capita from extreme heat of around 1.7% in the Midwest and Northeast whereas the OLS estimator projects losses to income per capita of around 1% in the Midwest and around 0.02% in the Northeast. The 95% confidence intervals for the ILS and OLS estimators do not overlap in these regions, illustrating the importance of using the ILS estimator. The projected effects on the South and West are smaller, with the ILS estimator projecting benefits from additional extreme heat in the South.

The second column under each estimator combines effects across all four weather bins. Total projected losses from climate change are typically larger than projected losses from additional extreme heat alone but in general comprise mostly losses from extreme heat.[40] The Midwest remains the region with the largest estimated losses. In particular, the ILS estimator projects reductions in output per capita in the Midwest of around 4%, but the effect is imprecisely estimated due to the imprecisely estimated effects of cold and cool weather on output per capita in that region. The ILS estimator projects statistically and economically significant reductions in income per capita in every region but the South. Those losses are around 2% in both the Midwest and Northeast and are around 0.6% in the West, with benefits in the South that are smaller than suggested by extreme heat alone and not statistically distinguishable from zero.

[40] Appendix E reports effects of climate change by temperature bin. Of note, additional cold days are beneficial, although the benefits are small relative to the costs of additional hot days. It is possible that the estimated benefits of cold days reflect correlation with reduced leisure or with omitted precipitation variables (e.g., the coldest days may have less snow). Because precipitation is difficult to project reliably in models of climate change, I here estimate effects of temperature inclusive of correlated changes in precipitation and project future climate impacts as if that correlation will be preserved in the future.

Table 3: Projected end-of-century percentage change in output and income per capita due to climate change.

	ILS		OLS	
	Days Above 80 degF	All Days	Days Above 80 degF	All Days
		Output p.c. (%)		
Midwest	-2.33	-4.28	-0.96	-1.25
	(-3.59,-1.47)	(-12.0,1.98)	(-1.28,-0.63)	(-1.65,-0.88)
Northeast	-0.26	-0.48	-0.051	-0.092
	(-6.20,0.34)	(-6.50,1.66)	(-0.11,-0.00077)	(-0.17,-0.012)
South	0.68	0.32	-0.32	-0.56
	(-0.067,1.74)	(-0.98,1.96)	(-0.59,-0.080)	(-0.81,-0.31)
West	0.19	-0.024	0.044	0.031
	(-0.025,0.48)	(-0.60,0.45)	(-0.043,0.15)	(-0.10,0.18)
		Income p.c. (%)		
Midwest	-1.65	-2.18	-1.03	-1.59
	(-1.99,-1.37)	(-2.69,-1.81)	(-1.19,-0.85)	(-1.82,-1.40)
Northeast	-1.67	-1.93	-0.017	-0.052
	(-293.0,-0.92)	(-293.1,-1.04)	(-0.038,0.0046)	(-0.080,-0.023)
South	0.60	0.33	-0.24	-0.44
	(0.19,1.20)	(-0.088,0.84)	(-0.36,-0.13)	(-0.57,-0.31)
West	-0.50	-0.55	-0.0082	-0.071
	(-60.5,-0.20)	(-60.5,-0.23)	(-0.054,0.044)	(-0.12,-0.011)

95% confidence intervals (in parentheses) bootstrapped from 1000 samples with resampling at the county level.
Projected temperature changes average 30 CMIP6 models, for RCP 4.5 in 2081–2100 relative to 1995–2014.
Changes in annual days above 80deg F: 12.1 in Midwest, 1.9 in Northeast, 21.0 in South, 2.6 in West.
ILS columns multiply projected temperature changes by indirect least squares estimates for total effects, from Proposition 7 and regression (21).
OLS columns multiply projected temperature changes by $\hat{\theta}_{rk}$ from regression (22).

The ILS projections are again typically more severe than—and significantly different from—the OLS projections.

Theory showed that these projected consequences of climate change are partially identified, before considering the possibility of ex-ante adaptation and fixed-cost adaptation. The positive estimates for $\hat{\phi}_{2r4}/\hat{\phi}_{1r4}$ in Table 2 and the industry-level analysis in Appendix C both suggest that, for the critical extreme heat variable, ex-post adaptation will be greater in the long run than the short run. Because ex-post adaptation estimates reported in Table 2 and Appendix E tend to have the same sign as the estimated total effects, the estimated effects of climate change are smaller than the true steady-state effects.[41] From the smaller end of the ILS estimator's 95% confidence interval, we can conclude that, barring ex-ante adaptation, reductions in income per capita are likely to be at least 1.8% in the Midwest, at least 1% in the Northeast, and at least 0.23% in the West. In each case, the lower bound projects larger losses than the OLS estimator would have, as the lower bound is either clearly more negative than or just on the edge of the OLS estimator's 95% confidence interval.

7 Discussion

I have explored the limits of our ability to estimate the long-run effects of climate change purely from short-run, panel variation in weather that is clearly exogenous, without postulating variation in climate either cross-sectionally or over time and without postulating that we can observe agents' adaptation choices. I have formally shown that we can bound long-run effects by using a new indirect least squares estimator. In an application to U.S. counties' output and income, I have empirically shown that the new estimator can generate significantly different conclusions from conventional estimators that are not grounded in theory. Future work should apply these new methods to other settings, including ones in which observable forecasts enable estimates of ex-ante adaptation.

I have highlighted the dynamic differences between transient weather shocks and permanent shifts in climate. Weather shocks and climate change also differ in their spatial structure and thus in their general equilibrium implications. Future work should explore how to credibly conduct inference about climate change from weather in these other dimensions (see Desmet and Rossi-Hansberg, 2024). In addition, I have followed prior empirical literature in estimating the effects of changing one stationary climate to another. Future work should consider the process of changing the climate. Finally, by imposing stronger assumptions on the decision-making environment and calibrating parameters to target the long-run costs implied by the methods presented here, future work could simulate counterfactual climate trajectories and

[41] And because ex-post adaptation and direct effects work in the same direction in this application, we know that the steady-state effects of climate change are at least as large as the estimated direct effects, even without knowing which way estimated ex-post adaptation is biased.

thereby estimate the costs of transitioning from one climate to another.

References

Aldy, Joseph E. and Richard Zeckhauser (2020) "Three prongs for prudent climate policy," *Southern Economic Journal*, Vol. 87, No. 1, pp. 3–29.

Aragón, Fernando M., Francisco Oteiza, and Juan Pablo Rud (2021) "Climate change and agriculture: Subsistence farmers' response to extreme heat," *American Economic Journal: Economic Policy*, Vol. 13, No. 1, pp. 1–35.

Auffhammer, Maximilian (2018) "Quantifying economic damages from climate change," *Journal of Economic Perspectives*, Vol. 32, No. 4, pp. 33–52.

——— (2022) "Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption," *Journal of Environmental Economics and Management*, Vol. 114, p. 102669.

Bakkensen, Laura and Lint Barrage (2024) "Climate shocks, cyclones, and economic growth: Bridging the micro-macro gap," *The Economic Journal*, Vol. Conditionally accepted.

Barker, David (2022) "Temperature and U.S. Economic Growth: Comment on Colacito, Hoffmann, and Phan," *Econ Journal Watch*, Vol. 19, No. 2, pp. 176–189.

Bento, Antonio, Noah S Miller, Mehreen Mookerjee, and Edson R Severnini (2023) "A unifying approach to measuring climate change impacts and adaptation," *Journal of Environmental Economics and Management*, Vol. 121, p. 102843.

Blanc, Elodie and Wolfram Schlenker (2017) "The use of panel models in assessments of climate impacts on agriculture," *Review of Environmental Economics and Policy*, Vol. 11, No. 2, pp. 258–279.

Burke, Marshall and Kyle Emerick (2016) "Adaptation to climate change: Evidence from US agriculture," *American Economic Journal: Economic Policy*, Vol. 8, No. 3, pp. 106–140.

Carleton, Tamma A. and Solomon M. Hsiang (2016) "Social and economic impacts of climate," *Science*, Vol. 353, No. 6304, p. aad9837.

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang (2022) "Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits," *The Quarterly Journal of Economics*, Vol. 137, No. 4, pp. 2037–2105.

Carter, Colin, Xiaomeng Cui, Dalia Ghanem, and Pierre Mérel (2018) "Identifying the economic impacts of climate change on agriculture," *Annual Review of Resource Economics*, Vol. 10, No. 1, pp. 361–380.

Chetty, Raj (2009) "Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods," *Annual Review of Economics*, Vol. 1, No. 1, pp. 451–488.

Colacito, Riccardo, Bridget Hoffmann, and Toan Phan (2019) "Temperature and growth: A panel analysis of the United States," *Journal of Money, Credit and Banking*, Vol. 51, No. 2-3, pp. 313–368.

Cui, Xiaomeng (2020) "Beyond yield response: Weather shocks and crop abandonment," *Journal of the Association of Environmental and Resource Economists*, Vol. 7, No. 5, pp. 901–932, Publisher: The University of Chicago Press.

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*, Vol. 4, No. 3, pp. 66–95.

——— (2014) "What do we learn from the weather? The new climate-economy literature," *Journal of Economic Literature*, Vol. 52, No. 3, pp. 740–798.

Deryugina, Tatyana and Solomon Hsiang (2017) "The marginal product of climate," Working Paper 24072, National Bureau of Economic Research.

Deschênes, Olivier and Michael Greenstone (2007) "The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather," *American Economic Review*, Vol. 97, No. 1, pp. 354–385.

——— (2012) "The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Reply," *American Economic Review*, Vol. 102, No. 7, pp. 3761–3773.

Desmet, Klaus and Esteban Rossi-Hansberg (2024) "Climate change economics over time and space," March.

Fankhauser, Samuel, Joel B. Smith, and Richard S. J. Tol (1999) "Weathering climate change: some simple rules to guide adaptation decisions," *Ecological Economics*, Vol. 30, No. 1, pp. 67–78.

Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker (2012) "The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather: Comment," *American Economic Review*, Vol. 102, No. 7, pp. 3749–3760.

Fried, Stephie (2022) "Seawalls and stilts: A quantitative macro study of climate adaptation," *The Review of Economic Studies*, Vol. 89, No. 6, pp. 3303–3344.

Graff Zivin, Joshua and Matthew Neidell (2009) "Days of haze: Environmental information disclosure and intertemporal avoidance behavior," *Journal of Environmental Economics and Management*, Vol. 58, No. 2, pp. 119–128.

Guo, Christopher and Christopher Costello (2013) "The value of adaption: Climate change and timberland management," *Journal of Environmental Economics and Management*, Vol. 65, No. 3, pp. 452–468.

Hamermesh, Daniel S. (1995) "Comment," *ILR Review*, Vol. 48, No. 4, pp. 835–838.

Hausfather, Zeke and Glen P. Peters (2020) "Emissions—the 'business as usual' story is misleading," *Nature*, Vol. 577, No. 7792, pp. 618–620.

Heal, Geoffrey and Jisung Park (2016) "Temperature stress and the direct impact of climate change: A review of an emerging literature," *Review of Environmental Economics and Policy*, Vol. 10, No. 2, pp. 347–362.

Heckman, James J. (2010) "Building bridges between structural and program evaluation approaches to evaluating policy," *Journal of Economic Literature*, Vol. 48, No. 2, pp. 356–398.

Hennessy, Christopher A. and Ilya A. Strebulaev (2020) "Beyond random assignment: Credible inference of causal effects in dynamic economies," *Journal of Finance*, Vol. 75, No. 2, pp. 825–866.

Hsiang, Solomon M. (2016) "Climate econometrics," *Annual Review of Resource Economics*, Vol. 8, pp. 43–75.

Huntingford, Chris, Philip D. Jones, Valerie N. Livina, Timothy M. Lenton, and Peter M. Cox (2013) "No increase in global temperature variability despite changing regional patterns," *Nature*, Vol. 500, No. 7462, pp. 327–330.

Jagnani, Maulik, Christopher B. Barrett, Yanyan Liu, and Liangzhi You (2021) "Within-season producer response to warmer temperatures: Defensive investments by Kenyan farmers," *The Economic Journal*, Vol. 131, No. 633, pp. 392–419.

Judd, Kenneth L. (1996) "Approximation, perturbation, and projection methods in economic analysis," in Hans M. Amman, David A. Kendrick, and John Rust eds. *Handbook of Computational Economics*, Vol. 1: Elsevier, pp. 509–585.

Keane, Michael P. and Kenneth I. Wolpin (2002) "Estimating welfare effects consistent with forward-looking behavior. Part I: Lessons from a simulation exercise," *The Journal of Human Resources*, Vol. 37, No. 3, pp. 570–599.

Kelly, David L., Charles D. Kolstad, and Glenn T. Mitchell (2005) "Adjustment costs from environmental change," *Journal of Environmental Economics and Management*, Vol. 50, No. 3, pp. 468–495.

Kolstad, Charles D. and Frances C. Moore (2020) "Estimating the economic impacts of climate change using weather observations," *Review of Environmental Economics and Policy*, Vol. 14, No. 1, pp. 1–24.

Kruttli, Mathias, Brigitte Roth Tran, and Sumudu W. Watugala (2019) "Pricing Poseidon: Extreme weather uncertainty and firm return dynamics," Finance and Economics Discussion Series 2019-054, Board of Governors of the Federal Reserve System.

Leduc, Sylvain and Daniel J. Wilson (2023) "Climate change and the geography of the U.S. economy," Working Paper 2023-17, Federal Reserve Bank of San Francisco.

Lemoine, Derek and Sarah Kapnick (2016) "A top-down approach to projecting market impacts of climate change," *Nature Climate Change*, Vol. 6, No. 1, pp. 51–55.

——— (2024) "Financial markets value skillful forecasts of seasonal climate," *Nature Communications*, Vol. forthcoming.

Lucas, Robert E., Jr. (1972) "Econometric testing of the natural rate hypothesis," in Otto Eckstein ed. *The Econometrics of Price Determination*: Board of Governors of the Federal Reserve System, pp. 50–59.

Massetti, Emanuele and Robert Mendelsohn (2018) "Measuring climate adaptation: Methods and evidence," *Review of Environmental Economics and Policy*, Vol. 12, No. 2, pp. 324–341.

Mendelsohn, Robert (2000) "Efficient adaptation to climate change," *Climatic Change*, Vol. 45, No. 3-4, pp. 583–600.

——— (2019) "An examination of recent revealed preference valuation methods and results," *Review of Environmental Economics and Policy*, Vol. 13, No. 2, pp. 267–282.

Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw (1994) "The impact of global warming on agriculture: A Ricardian analysis," *The American Economic Review*, Vol. 84, No. 4, pp. 753–771.

Mérel, Pierre and Matthew Gammans (2021) "Climate econometrics: Can the panel approach account for long-run adaptation?" *American Journal of Agricultural Economics*, Vol. 103, No. 4, pp. 1207–1238.

Mohaddes, Kamiar, Ryan N C Ng, M Hashem Pesaran, Mehdi Raissi, and Jui-Chung Yang (2023) "Climate change and economic activity: Evidence from US states," *Oxford Open Economics*, Vol. 2, p. odac010.

Newell, Richard G., Brian C. Prest, and Steven E. Sexton (2021) "The GDP-temperature relationship: Implications for climate change damages," *Journal of Environmental Economics and Management*, Vol. 108, p. 102445.

Nordhaus, William D. (2006) "Geography and macroeconomics: New data and new findings," *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 103, No. 10, pp. 3510–3517.

Quiggin, John and John Horowitz (2003) "Costs of adjustment to climate change," *Australian Journal of Agricultural and Resource Economics*, Vol. 47, No. 4, pp. 429–446.

Quiggin, John and John K. Horowitz (1999) "The impact of global warming on agriculture: A Ricardian analysis: Comment," *The American Economic Review*, Vol. 89, No. 4, pp. 1044–1045.

Rosenzweig, Mark and Christopher R. Udry (2013) "Forecasting profitability," Working Paper 19334, National Bureau of Economic Research.

Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher (2005) "Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach," *American Economic Review*, Vol. 95, No. 1, pp. 395–406.

Schlenker, Wolfram and Michael J. Roberts (2006) "Nonlinear effects of weather on corn yields," *Applied Economic Perspectives and Policy*, Vol. 28, No. 3, pp. 391–398.

——— (2009) "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change," *Proceedings of the National Academy of Sciences*, Vol. 106, No. 37, pp. 15594–15598.

Severen, Christopher, Christopher Costello, and Olivier Deschênes (2018) "A forward-looking Ricardian approach: Do land markets capitalize climate change forecasts?" *Journal of Environmental Economics and Management*, Vol. 89, pp. 235–254.

Shrader, Jeffrey (2020) "Expectations and adaptation to environmental risks," working paper.

Shrader, Jeffrey G., Laura Bakkensen, and Derek Lemoine (2023) "Fatal errors: The mortality value of accurate weather forecasts," Working Paper 31361, National Bureau of Economic Research.

Sorkin, Isaac (2015) "Are there long-run effects of the minimum wage?" *Review of Economic Dynamics*, Vol. 18, No. 2, pp. 306–333.

Thrasher, Bridget, Weile Wang, Andrew Michaelis, Forrest Melton, Tsengdar Lee, and Ramakrishna Nemani (2022) "NASA Global Daily Downscaled Projections, CMIP6," *Scientific Data*, Vol. 9, No. 1, p. 262, Number: 1 Publisher: Nature Publishing Group.

Tinbergen, J. (1930) "Bestimmung und deutung von angebotskurven ein beispiel," *Zeitschrift für Nationalökonomie*, Vol. 1, No. 5, pp. 669–679.

——— (1995) "Determination and interpretation of supply curves: An example [Bestimmung und deutung von angebotskurven. ein beispiel] (Zeitschrift fur Nationalokonomie, vol. 1, 1930, pp. 669–79)," in David F. Hendry and Mary S. Morgan eds. *The Foundations of Econometric Analysis*, Cambridge: Cambridge University Press, pp. 233–246.

Weyl, E. Glen (2019) "Price theory," *Journal of Economic Literature*, Vol. 57, No. 2, pp. 329–384.