

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 2021

I thank numerous seminar participants for valuable comments. I thank Wolfram Schlenker for kindly providing data, Emily Becker for reanalyzing her data, and Olivier Deschênes and Ryan Kellogg for helpful conversations. 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Derek Lemoine. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Estimating the Consequences of Climate Change from Variation in Weather
Derek Lemoine
NBER Working Paper No. 25008
September 2018, Revised May 2021
JEL No. D84,H43,Q54

ABSTRACT

I formally relate the consequences of climate change to the panel variation in weather extensively explored by recent empirical literature. I show that short-run responses to weather shocks differ from long-run responses to climate change when payoffs depend on a capital or resource stock. I develop a new indirect least squares estimator that bounds long-run climate impacts from short-run responses to weather. Applying this new method, I find that an additional 2°C of global warming would eliminate profits from the average acre of current farmland in the eastern U.S.

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

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

1 Introduction

A pressing research agenda seeks to estimate the economic costs of climate change. Ignorance of these costs has hampered policy. 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, leading to concerns about omitted variables bias.¹

Intriguingly, though, the same location does experience different weather at different times. Stimulated by Deschênes and Greenstone (2007), 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 with scientists' predictions about how climate change will alter the distribution of weather. But it has been unclear whether extrapolating estimated effects of weather is truly informative about climate change impacts.³

This paper formally relates the effects of climate change to the effects of weather shocks. I focus on the different dynamic structure of weather shocks and climate change: weather shocks are transient by construction, whereas climate change permanently alters the distribution of weather. I show that estimating the 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).

The best possible weather regressions suffer from three biases when used to project climate change impacts. All three biases derive from the adaptation channel. First, some types of long-lived infrastructure will respond to climate change but are fixed in the data. This infrastructure causes bias only if its adjustments affect the shorter-run adaptation decisions that do vary in the data. A second bias arises from using transient shocks to weather forecasts to estimate ex-ante adaptation. I show that reactions to weather forecasts identify a combination of the ex-ante adaptation relevant to climate change and preparation for transient changes in ex-post adaptation. The latter is not relevant to the effects of systematically altering average forecasts across periods, as with climate change that affects agents' expectations in every period.

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

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

³For 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.”

The third bias reflects the difference between experiencing a transient weather shock and living with altered weather period after period, as after a change in climate. Actions are intertemporal complements (substitutes) if actions in one period increase (decrease) 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. In the former case, estimates derived from short-run weather variation 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?⁵ I develop a new indirect least squares estimator of climate impacts (Tinbergen, 1930, 1995). I show that we can in fact partially identify the long-run effects of climate change in a fairly general setting, even without observing all the actions agents and firms could choose, without observing all the capital and resource stocks that they interact with, and without assuming functional forms. I first express climate change impacts in terms of theoretical primitives. I then derive what reduced-form weather coefficients estimate within this general model. Finally, I invert the system of reduced-form coefficients to recover combinations of theoretical primitives that I insert into the expression for climate impacts. The identification is purely reduced-form, as the only estimation is a fixed effects regression that relies on panel variation in weather, but both the specification of the regression and the calculations with its estimates derive from theory.⁶

⁴Both types of stories exist in the literature (see Auffhammer, 2018b). 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.

⁵Recent literature has sought to work around concerns about the relevance of short-run variation in weather by estimating how the effect of weather varies cross-sectionally with a location's climate (e.g., Auffhammer, 2018a; Carleton et al., 2020) or by using "long difference" estimators (e.g., Dell et al., 2012; Burke and Emerick, 2016). The former approach forsakes the clean identification of panel variation; I instead explore the limit of what researchers can learn from purely panel variation in weather. Appendix A analyzes long difference estimators, showing that they inherit the biases suffered by standard weather regressions. Finally, other work uses quasi-random spatial variation in water supplies to estimate long-run adaptation (e.g., Hornbeck and Keskin, 2014; Blakeslee et al., 2020; Hagerty, 2020), but similar variation will not be available for many environmental variables affected by climate change.

⁶Critically, this calculation does not require the specification of structural parameters or even of functional forms. 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

I show that this indirect least squares estimator eliminates the bias induced by preparation for transient changes in ex-post adaptation. It also signs the wedge between short-run and long-run adaptation induced by the transience of weather shocks. Because this estimator decomposes climate impacts into direct effects of weather and adaptation channels, I use that sign to bound the effects of climate change. The remaining bias is the possibility that some long-lived infrastructure could adjust on timescales not observed in the data and thereby alter the adaptation responses that are recovered from the data. Such bias would matter only if it were in a direction that would violate the estimated bound.

I demonstrate this new method with an updated version of a seminal analysis of climate and agriculture (Deschênes and Greenstone, 2007). Conventional reduced-form calculations suggest that 2°C of global warming would reduce profits from the average acre in the eastern U.S. by around 42%, driven by changes in extreme heat. However, the model primitives recovered by indirect least squares reject the assumptions that I show are necessary for the 42% calculation to be valid.

My new indirect least squares estimates imply that 2°C of global warming would largely eliminate profits from the average acre of farmland in the eastern U.S. The critical difference is the effect of common heat (“growing degree days”). The conventional reduced-form regression suggests that agricultural profits benefit from additional growing degree days, but I show that this estimate entwines direct effects of heat with the short-run effects of ex-post adaptation. Using the effects of lagged weather to disentangle these, I find that ex-post adaptation provides short-run benefits. For dynamically optimizing agents, these nonzero marginal benefits in the short run reflect tradeoffs with long-run costs, as when adapting through increased use of a scarce natural resource.⁷ Ex-post adaptation therefore reduces the near-term costs of climate change but increases the long-run costs of climate change. Moreover, once we clean the coefficient on contemporary weather of the short-run benefits of ex-post adaptation, the direct effects of additional growing degree days are harmful. Whereas the reduced-form analysis suggested that additional growing degree days mitigate the costs of climate change, we see that these in fact increase the long-run costs of climate change both through direct effects and through adaptation that imposes dynamic tradeoffs.

If the combination of ex-post and ex-ante adaptation increases the long-run costs of climate change, then I can recover a lower bound on the cost of climate change from the estimated direct effects of climate change. The indirect least squares estimator shows that actions are intertemporal substitutes, which is again consistent

knowledge of individual structural parameters.” It is also related to sufficient statistics approaches (see Chetty, 2009) and to price theory (see Weyl, 2019).

⁷For example, Blakeslee et al. (2020) show that Indian households adapt to water scarcity by accumulating debt and removing children from school, both of which impose long-run costs. And Aragón et al. (2021) show that Peruvian farmers increase area planted in response to high temperatures, which they speculate could reduce future land productivity.

with adaptation depending on scarce resources. In that case, short-run adaptation is greater than long-run adaptation and I can recover an upper bound on the costs of climate change by summing the estimated costs of adaptation and the estimated direct effects. However, identifying ex-ante adaptation is especially challenging in this application because I do not observe forecasts of weather several months to a year ahead. I instead proxy for forecasts with the lead of weather and calibrate the bias induced by this proxy. The resulting estimates are too imprecise to yield a useful upper bound on costs, but I do show that total adaptation does increase costs under plausible calibrations. Because the estimated direct effects of climate change then do provide a lower bound on the costs of climate change, I conclude that 2°C of global warming would largely eliminate profits from the average acre of farmland.

There has been remarkably little prior formal analysis of the economic link between weather and climate, despite the importance of empirically estimating the costs of climate change and the sharpness of informal debates around the relevance of the burgeoning empirical literature to climate change. The primary exceptions are Hsiang (2016) and Deryugina and Hsiang (2017). They argue that the simplest weather regression exactly identifies the effect of climate on payoffs. In their setting, outcomes and actions depend only on the distribution of weather (i.e., only on the climate), not on the weather realized from this distribution. In Section 3, this formulation will emerge as a special case of the present setting. I show that the simplest weather regressions do then recover the effect of climate, and I show that this optimistic result survives allowing actions to respond to realized weather. However, I also show that it does not survive allowing actions to be dynamically linked. If either current actions can protect against future weather or payoffs depend on a capital or resource stock inherited from earlier periods, then the short-run effects of transient weather shocks are no longer identical to the long-run effects of climate.⁸

The challenge of attempting to estimate long-run effects from short-run variation is a common one in empirical economics. The present analysis and methods could inform approaches in other fields. 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

⁸Shrader (2020) shows that forecasts are valuable for disentangling ex-ante adaptation from ex-post adaptation and the direct effects of weather. I show an analogous result when seeking to infer effects of climate change. I also show that estimating the effects of lagged weather allows ex-post adaptation to be disentangled from direct effects, that using multiple forecast horizons can eliminate one source of bias when extrapolating to climate change, and that leads of weather can imperfectly proxy for unobserved forecasts.

⁹Three 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

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. 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 develops the new method of estimating climate impacts and applies it to U.S. agriculture. The final section describes potential extensions. The appendix contains empirical details, additional analysis, proofs, and robustness checks.

2 Setting

In each period t , agents receive payoffs $\pi(w_t, A_t, S_t; K)$, with π bounded.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 $\pi_{SS} < 0$ (except in Section 3, where $\pi_S = 0$). 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. If $g > 0$, the time $t + 1$ stock depends on all past actions, as with investments in a capital stock that depreciates at rate $1 - g$.

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.¹² 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 cropping practices 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

and panel variation when estimating the effects of welfare benefits. These tradeoffs are similar to those that motivate the present paper.

¹⁰I refer to “agents” and “actions”, but one can instead think of firms choosing quantities, with weather affecting either prices or the production function. The assumption of boundedness is a technical condition that ensures optimal policy is single-valued (used in Appendix E.3).

¹¹I 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.

¹²Exogenous groundwater recharge is consistent with a constant in $h(A_t)$.

actions. As $|h' \pi_{AS}|$ becomes 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. We can interpret weather as realized temperature and climate as a location's long-run average temperature. 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: $f_{2,t-2} = C + \zeta \epsilon_{2,t-2}$.¹³ 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 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 $w_t = f_{1,t-1} + \zeta \epsilon_{0,t}$. Reflecting rationality of beliefs, the random variables are mean-zero and serially uncorrelated. Ordering the $\epsilon_{i,t}$ by i , they have covariance matrix Σ at any time t . 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.¹⁴

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

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

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:

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

Agents also choose long-lived infrastructure K . This represents capital-intensive adaptation that takes years to construct, such as irrigation canals or sea walls (see

¹³Implicitly, $f_{k,t} = C$ 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.

¹⁴Consistent 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 $\zeta^2 \text{trace}(\Sigma)$, 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.

Aldy and Zeckhauser, 2020). This infrastructure is fixed over the period of analysis; agents cannot adapt it to short-run weather outcomes or forecasts. This is the only kind of action analyzed in previous work that formally relates climate change to weather variation (Hsiang, 2016; Deryugina and Hsiang, 2017). The agent chooses K to maximize long-run payoffs under expected outcomes:¹⁵

$$\max_K \left\{ \lim_{t \rightarrow \infty} \pi(E_0[w_t], E_0[A_t], E_0[S_t]; K) \right\}.$$

Together, the decision variables A_t and K bracket the many types of actions actual agents may take, which fall on a spectrum between the immediate consequences of changing A_t and the purely long-run consequences of changing K .

The setting is meant to be fairly general. To fix ideas, consider a few examples pertinent to previous literature. In an agricultural application, actions could be planting decisions, the stock could be water supplies or machinery, and long-lived infrastructure could be irrigation canals or available crop varieties.¹⁶ In a flooding application, the 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, the “stock” (i.e., the inherited state) could be one’s current location and the actions could be the choice of future location. That choice fits the formal framework if it depends on the current location’s present and expected weather and on the alternate location’s average weather. In a health application, individuals may organize their outdoor activities around weather forecasts in order to maximize utility net of health and mortality risks but find this ability restricted as more days go by and the stock of postponed activities accumulates (as in Graff Zivin and Neidell, 2009). In an innovation application, the stock could be existing patents on adaptation technologies, weather could affect the demand for these patents, and the action could be investing in research. And in a labor or energy application, weather could affect labor productivity or residential comfort, the stock could be air conditioning equipment, and the action could be investing in air conditioning.

I assume the following 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 E.3):

$$[h'(A_t)\pi_{AS}]^2 < [h'(A_t)]^2\pi_{SS} \left[\pi_{AA} - \frac{h''(A_t)}{h'(A_t)}\pi_A \right]. \quad (1)$$

Observe that inequality (1) and $\pi_{SS} < 0$ imply

$$\pi_{AA} - \frac{h''(A_t)}{h'(A_t)}\pi_A < 0. \quad (2)$$

¹⁵This form of the infrastructure decision problem does not sacrifice the qualitative insight of maximizing expected payoffs but does simplify exposition.

¹⁶Recent 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.

The next two conditions ensure that a steady state exists in a deterministic system with $\zeta = 0$ (Appendix E.4):

$$\lim_{A_t \rightarrow -\infty} -(1 - \beta g)\pi_A(C, A_t, \cdot; K) - \beta h'(\bar{A})\pi_S(C, A_t, \cdot; K) < 0, \quad (3)$$

$$\lim_{A_t \rightarrow \infty} -(1 - \beta g)\pi_A(C, A_t, \cdot; K) - \beta h'(\bar{A})\pi_S(C, A_t, \cdot; K) > 0. \quad (4)$$

The final condition ensures that the expression for expected optimal actions converges (Lemma 2 in Appendix E.6):

$$h'(A_t)\pi_{AS} \in \left(-\frac{[1 + 2g(1 + \beta) + 3\beta g^2] \left[-\pi_{AA} + \frac{h''(A_t)}{h'(A_t)}\pi_A \right] - \beta[h'(A_t)]^2\pi_{SS}}{1 + \beta + 2\beta g}, \frac{[1 - 2g(1 + \beta) + 3\beta g^2] \left[-\pi_{AA} + \frac{h''(A_t)}{h'(A_t)}\pi_A \right] - \beta[h'(A_t)]^2\pi_{SS}}{1 + \beta - 2\beta g} \right). \quad (5)$$

The interval includes zero. This condition therefore permits both intertemporal complementarity and intertemporal substitutability but limits the degree of either.

The analysis approximates the solution to the full, stochastic model around the steady state of the deterministic model, which has $\zeta = 0$ (Judd, 1996). In order to ensure an adequate approximation, I will often impose at least one of the following assumptions:

Assumption 1. ζ^2 is small.

Assumption 2. π is quadratic.

Assumption 3. The $\epsilon_{i,t}$ are jointly normally distributed.

Either of the first two assumptions will limit the consequences of stochasticity for optimal policy, whether by limiting the variance of weather outcomes (Assumption 1) or by making the policy function independent of that variance (Assumption 2).¹⁷ And using either the first or the third assumption will eliminate covariances of certain higher-order terms.

I am interested in empirical researchers' ability to estimate the consequences of altering C from observable responses to panel variation in w_t , $f_{1,t}$, and $f_{2,t}$. I assume that empirical researchers observe 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 S .¹⁸

¹⁷When applying Assumption 2, the *chosen policy* is indeed affected by the variance of weather (through realized weather) even though the *policy rule* is independent of that variance.

¹⁸Omitted variables bias affects the analysis below when regressions do not control for variables (such as forecasts and actions) that are defined within the theoretical model. I do not explicitly model the further unobservable characteristics that motivate fixed effects specifications, as I am here

Finally, it is important to be clear about the treatment effect of interest. I will 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. The climate change treatment is consistent with the dominant exercise in the empirical literature to date, which typically calculates the effect of replacing today’s distribution of weather with a distribution projected to hold by the end of the century. Following this literature, I will not study how the transition from one climate to another interacts with agents’ decisions¹⁹ or study how expectations of a future change in climate affect agents today.²⁰ 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 $\pi_A(w_t, A_t^*, S_t; K) = 0$ and independent of all other periods’ actions.²¹

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

not interested in whether regression coefficients are identified in the applied microeconomics sense of consistently estimating weather impacts but in whether climate change impacts are identified by weather in the econometric sense of recoverable from data. The applied microeconomics notion of identification motivated the literature’s shift from cross-sectional to panel variation (see Dell et al., 2014; Auffhammer, 2018b). This paper explores only panel variation.

¹⁹Kelly et al. (2005) frame the cost of learning as an adjustment cost. Quiggin and Horowitz (1999, 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.

²⁰Severen 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).

²¹Using terminology defined below, there is no ex-ante adaptation and ex-post adaptation is not affected by previous decisions. I recover the setting of Hsiang (2016) and Deryugina and Hsiang (2017) if I further eliminate the choice of A_t (and also 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.

either Assumption 1 or 2,

$$\frac{dE_0[\pi_t]}{dC} = \bar{\pi}_w + \underbrace{\bar{\pi}_A \frac{d\bar{A}}{dC}}_{=0} + \underbrace{\bar{\pi}_K \frac{d\bar{K}}{dC}}_{=0} = \bar{\pi}_w \quad (6)$$

for $t > 2$. 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).²²

Consider the following regression

$$\pi_{jt} = \alpha_j + \theta w_{jt} + \eta_{jt}, \quad (7)$$

where α_j is a fixed effect for unit j and η_{jt} is an error term. Use a hat to denote the probability limit of each estimator. By standard results,

$$\hat{\theta} = \frac{Cov[\pi_{jt}, w_{jt} - C]}{Var[w_{jt} - C]}.$$

We now have:

Proposition 1. *Let Assumption 1 hold, or let Assumptions 2 and 3 hold. If $\pi_S = 0$, then $\hat{\theta} = \bar{\pi}_w$.*

Proof. See Appendix E.2. □

Therefore, from equation (6),

$$\frac{dE_0[\pi_t]}{dC} = \hat{\theta}$$

for $t > 2$. The simplest weather regression recovers the average marginal effect of weather and thus recovers the long-run effects of climate, as claimed by Hsiang (2016) and Deryugina and Hsiang (2017).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 potentially 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.,

²²Guo 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.

²³Much literature regresses outcomes other than payoffs on weather. It is easy to show that the coefficient on weather in a regression with actions as the dependent variable recovers the long-run effect of climate on actions.

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 hurricane forecasts (Krutli et al., 2019), people respond to environmental warnings (Neidell, 2009), and people value weather forecasts (Lazo et al., 2009). We next turn to the full setting to see how far the optimism engendered by the present specialization has to run.

4 How Dynamics Complicate Reduced-Form Approaches to Estimating Climate Impacts

With $\pi_S \neq 0$, agents must account for future consequences when choosing their actions. Appendix E.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 $\bar{\pi}$. And assume henceforth that agents are not too far from the steady state at time 0 (i.e., that $(S_0 - \bar{S})^2$ is not too large).

I first define the true effect of climate. I then describe 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 weather shocks.

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 E.7 shows that, if either Assumption 1 or 2 holds,

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{dE_0[\pi_t]}{dC} &= \bar{\pi}_w + \bar{\pi}_A \frac{d\bar{A}}{dC} + \bar{\pi}_S \frac{d\bar{S}}{dC} + \overbrace{\bar{\pi}_K \frac{d\bar{K}}{dC}}^{=0} \\ &= \bar{\pi}_w + \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g} \right] \frac{d\bar{A}}{dC}. \end{aligned} \quad (8)$$

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 have first-order consequences for average payoffs, both directly and through their effects on the stock.

Why do adaptation responses suddenly 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 $\pi_A = 0$.

Optimal actions satisfy the Euler equation, derived in Appendix E.5:

$$-\pi_A(w_t, A_t, S_t; K) = \beta h'(A_t) E_t \left[\pi_S(w_{t+1}, A_{t+1}, S_{t+1}; K) + g \frac{-\pi_A(w_{t+1}, A_{t+1}, S_{t+1}; K)}{h'(A_{t+1})} \right]. \quad (9)$$

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 include the direct 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 (as $\beta \rightarrow 0$) or as the stock becomes independent of actions (as $h' \rightarrow 0$).

We will therefore need to estimate how climate affects actions around the deterministic steady state \bar{A} if we are to recover the effect of climate on average payoffs. Appendix E.8 shows that

$$\frac{d\bar{A}}{dC} \propto \underbrace{\bar{\pi}_{wA}}_{\text{ex-post adaptation}} + \beta \underbrace{[h'(\bar{A})\bar{\pi}_{wS} - g\bar{\pi}_{wA}]}_{\text{ex-ante adaptation}} + \underbrace{[(1 - \beta g)\bar{\pi}_{AK} + \beta h'(\bar{A})\bar{\pi}_{SK}]}_{\text{interactions with long-lived infrastructure}} \frac{dK}{dC}. \quad (10)$$

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; Mendelsohn, 2000). 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. Ex-post adaptation can also reflect a firm's production responses 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). 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 payoffs when optimized ($\pi_K = 0$), but they do indirectly affect payoffs when the marginal benefit of either short-run actions or the stock 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.

where

$$\begin{aligned} \Psi &\triangleq h'(\bar{A})\bar{\pi}_{AS} + g \left(-\bar{\pi}_{AA} + \frac{h''(\bar{A})}{h'(\bar{A})}\bar{\pi}_A \right) \\ &\propto \left. \frac{dA_t}{dA_{t+1}} \right|_{w_t=f_1, t=f_2, t=C}. \end{aligned} \quad (13)$$

As should be expected, the coefficients on forecasts go to zero as agents become myopic. For forward-looking agents, three terms in equation (12) 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 cut back on current actions. This effect vanishes as $g \rightarrow 0$ 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 (10).

Third, agents anticipate how today's choices impose historical restraints on future choices and 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 (13) 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 $g \rightarrow 0$. The preparatory action term in equation (12) was absent from equation (10), a point 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 also affects agents' expectations of future weather, manifested as systematically higher forecasts. This channel will make it important to estimate the effects of forecasts.

I assume that the empirical researcher can observe payoffs (e.g., profits) and weather variables.²⁵ Importantly, I do not assume that the empirical researcher observes all of the actions that agents take or the level of the stock.²⁶ Consider the following distributed lag regression with fixed effects:

$$\pi_{jt} = \alpha_j + \sum_{i=0}^I \Lambda_i w_{j(t-i)} + \sum_{i=0}^I \lambda_i f_{j1,(t-i)} + \sum_{i=0}^I \gamma_i f_{j2,(t-i)} + \eta_{jt}, \quad (14)$$

where I again label units as 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, I use a hat to denote the probability limit of each coefficient.

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

Proposition 2. *Let Assumption 1 hold, or let Assumptions 2 and 3 hold. Then:*

$$\lim_{I \rightarrow \infty} \sum_{i=0}^{I-2} [\hat{\Lambda}_i + \hat{\lambda}_i] = \bar{\pi}_w + \omega \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g} \right] \left(\left. \frac{d\bar{A}}{dC} \right|_{K \text{ fixed}} + \Omega \right), \quad (15)$$

where

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

and Ψ is defined in equation (13). If $\beta \Psi > 0$, then $\omega < 1$. If $\beta \Psi = 0$, then $\omega = 1$. If $\beta \Psi < 0$, then $\omega > 1$.

Proof. See Appendix E.9. □

The good news is that we come somewhat close to the true effect of climate derived in equation (8). In particular, we successfully capture the direct effect of weather and we capture effects proportional to ex-post and ex-ante adaptation.

However, we also see three wedges between the true effect in equation (8) and the estimated effect in (15). First, the change in steady state actions in equation (15) holds K fixed, but equation (10) showed that $d\bar{A}/dC$ generally depends on changes in K . 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

²⁵The analysis straightforwardly extends to the case where the empirical researcher instead observes actions, not payoffs, and seeks the effect of climate on actions.

²⁶In Section 5.1, I will analyze a case in which the empirical researcher does not observe forecasts.

²⁷The requirement that we estimate at least I lags even though we use only $I - 2$ lags avoids ancillary complications from omitted variables bias at the longest lags.

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 (10) 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 (if $dK/dC = 0$ in equation (10)) or if it does not directly interact with other decisions (if $\bar{\pi}_{AK} = \bar{\pi}_{SK} = 0$ in equation (10)).

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 (12). 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 quite 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. When actions are intertemporal complements, 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 much greater adjustment. We have $\omega < 1$ because responses to transient shocks overstate historical restraints in this case. Consistent with conjectures in Deschênes and Greenstone (2007), observable short-run adaptation is less than long-run adaptation.

In contrast, when actions are intertemporal substitutes, 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. We have $\omega > 1$ because responses to transient shocks can understate historical restraints in this case. Consistent with conjectures in Fisher et al. (2012) and Blanc and Schlenker (2017), observable short-run adaptation is greater than long-run

²⁸One could eliminate Ω by not using the forecast coefficients $\hat{\lambda}_i$, instead relying on $\lim_{I \rightarrow \infty} \sum_{i=0}^{I-2} \hat{\Lambda}_i$. However, this calculation would introduce a new bias, as it would miss all ex-ante adaptation terms in equation (10). 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 (15).

adaptation.²⁹

Proposition 2 described the results of estimating a model with infinite lags and summing the coefficients. The following corollary describes regressions with fewer lags:³⁰

Corollary 3. *Let $I' \geq 1$ and $I \geq I' + 2$. Also let Assumption 1 hold, or let Assumptions 2 and 3 hold. Then:*

$$\sum_{i=0}^{I'} [\hat{\Lambda}_i + \hat{\lambda}_i] = \bar{\pi}_w + \omega_{I'} \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g} \right] \left(\left. \frac{d\bar{A}}{dC} \right|_{K \text{ fixed}} + \Omega \right),$$

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 E.10. □

The number of summed lags only affects 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 $\omega > 1$, so that responses to weather shocks overestimate responses to long-run changes in climate. In this case, the bias $\omega_{I'}$ fluctuates around ω 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 potentially induced by K being fixed:

Corollary 4. *Let Assumption 1 hold, or let Assumptions 2 and 3 hold. Let $I' \geq 1$ and $I \geq I' + 2$. Then:*

$$\lim_{\beta \rightarrow 0} \lim_{I \rightarrow \infty} \sum_{i=0}^{I-2} \hat{\Lambda}_i = \lim_{g, \bar{\pi}_{AS} \rightarrow 0} \sum_{i=0}^{I'} [\hat{\Lambda}_i + \hat{\lambda}_i] = \bar{\pi}_w + \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g} \right] \left. \frac{d\bar{A}}{dC} \right|_{K \text{ fixed}}.$$

Proof. See Appendix E.11. □

²⁹If $g > 0$, we can have $\omega < 1$ even when actions are intertemporal substitutes. The reason is that an agent living in an altered climate would intentionally loosen historical restraints over time.

³⁰The 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.

³¹Appendix 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). While long difference estimators are motivated by the possibility that climate change has manifested itself over long timesteps, 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 (14) with $I \geq \Delta + 2$ lags.

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). This wedge also 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 altered climate.³² Therefore we recover the effect of climate when we estimate infinite lags as long as agents are myopic and long-lived infrastructure either is fixed or does not interact with shorter-run adaptation decisions.³³

Second, 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 (11), we lose the effects of past weather (see footnote 24). 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 (8), the contemporary effects identify $\bar{\pi}_w + \bar{\pi}_A \left. \frac{d\bar{A}}{dC} \right|_{K \text{ fixed}}$ and the lagged effects identify $h'(\bar{A})\bar{\pi}_S \left. \frac{d\bar{A}}{dC} \right|_{K \text{ fixed}}$. Therefore we recover the effect of climate when we estimate at least one lag of weather and forecasts as long as $\bar{\pi}_{AS}, g = 0$ and long-lived infrastructure either is fixed or does not interact with shorter-run adaptation decisions.

5 Estimating Climate Impacts Through Indirect Least Squares

We have thus far seen that we can 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 either is fixed or does not interact with shorter-run adaptation decisions, or if agents make

³²The bias introduced by $\omega_{I'}$ in Corollary 3 does not vanish as $\beta \rightarrow 0$: 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.

³³When agents are myopic, we do not need to estimate responses to forecasts (and should obtain $\hat{\lambda}_i = 0$ and $\hat{\gamma}_i = 0$ if we do).

decisions independently over time and long-lived infrastructure either is fixed or does not interact with shorter-run adaptation decisions. But while we have described the biases that arise when these conditions do not hold, we have not been able to sign that bias.

I now show how an indirect least squares estimator can bound climate impacts and disentangle direct effects from each type of adaptation. Importantly, this new approach maintains precisely the same credible identification from the reduced-form specifications. As we will see, these specifications suffice because we do not need to recover—or even specify—every underlying structural parameter in order to undertake the calculations implied by theory.

I first preview the plan of attack. Fixing K , substituting for $d\bar{A}/dC$, and substituting for $\bar{\pi}_S$ from the Euler equation (9), equation (8) becomes:

$$\lim_{t \rightarrow \infty} \left. \frac{dE_0[\pi_t]}{dC} \right|_{K \text{ fixed}} = \underbrace{\bar{\pi}_w}_{\text{direct effects}} - \underbrace{\frac{1-\beta}{\beta} \bar{\pi}_A \frac{\bar{\pi}_{wA}}{D}}_{\text{ex-post adaptation}} - \underbrace{\frac{1-\beta}{\beta} \bar{\pi}_A \frac{\beta[h'(\bar{A})\bar{\pi}_{wS} - g\bar{\pi}_{wA}]}{D}}_{\text{ex-ante adaptation}}, \quad (16)$$

with $D > 0$ itself a function of cross-partials.³⁴ We aim to recover each individual piece of this expression from the estimated coefficients of regression (14). In particular, we will separately recover the direct effects, the ex-post adaptation term, and the ex-ante adaptation term, eliminating Ω and signing the effect of the analogue of ω . We will bound the total effect of climate change by using these pieces and that sign.

The following proposition expresses several combinations of theoretical primitives as functions of the estimated coefficients from regression (14).

Proposition 5. *Let Assumption 1 hold, or let Assumptions 2 and 3 hold. For $I > 2$, we have:*

$$\bar{\pi}_A \frac{\Gamma}{h'(\bar{A}) \bar{\chi}} = \hat{\lambda}_0, \quad \bar{\pi}_A \frac{\bar{\pi}_{wA}}{h'(\bar{A}) \bar{\chi}} = \hat{\Lambda}_1 \frac{\hat{\lambda}_0}{\hat{\lambda}_1}, \quad \bar{\pi}_w = \hat{\Lambda}_0 - \hat{\Lambda}_1 \frac{\hat{\lambda}_0}{\hat{\lambda}_1}, \quad \Psi \propto \frac{\hat{\Lambda}_2}{\hat{\Lambda}_1}.$$

Proof. See Appendix E.12. □

The estimated coefficient $\hat{\lambda}_0$ on contemporary forecasts identifies terms related to ex-ante adaptation, and the estimated coefficient $\hat{\Lambda}_1$ on lagged weather identifies terms related to ex-post adaptation (with an adjustment identified by the ratio of forecast coefficients). 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. Subtracting off the ex-post adaptation term identified by the lag of

³⁴See equations (A-23) and (A-24) in Appendix E.8. Note that D absorbs the $1 - g$ in the denominator left after substituting the Euler equation.

weather isolates the direct effects. Finally, the ratio of coefficients on lagged weather identifies the preparatory action term Ψ and thus how actions are linked over time.

We use these results to calculate the overall effect of climate:

$$\begin{aligned}
\hat{\Lambda}_0 - \hat{\Lambda}_1 \frac{\hat{\lambda}_0}{\hat{\lambda}_1} - \frac{1-\beta}{\beta} \left[\hat{\Lambda}_1 \frac{\hat{\lambda}_0}{\hat{\lambda}_1} + \hat{\lambda}_0 \right] &= \bar{\pi}_w - \frac{1-\beta}{\beta} \bar{\pi}_A \left\{ \frac{\bar{\pi}_{wA}}{h'(\bar{A})\bar{\chi}} + \frac{\Gamma}{h'(\bar{A})\bar{\chi}} \right\} \\
&= \bar{\pi}_w - \frac{D}{h'(\bar{A})\bar{\chi}} \frac{1-\beta}{\beta} \bar{\pi}_A \left(\frac{\bar{\pi}_{wA}}{D} + \frac{\beta[h'(\bar{A})\bar{\pi}_{wS} - g\bar{\pi}_{wA}]}{D} + \frac{\Omega}{1-g} \right) \\
&= \bar{\pi}_w + \frac{D}{h'(\bar{A})\bar{\chi}} \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g} \right] \left(\frac{d\bar{A}}{dC} \Big|_{K \text{ fixed}} + \Omega \right). \tag{17}
\end{aligned}$$

We calculate the left-hand side of the first line using the estimated coefficients and a calibrated value for β . The right-hand side of the first line uses Proposition 5 to express this calculation in terms of model primitives. The second line substitutes for Γ . Substituting $d\bar{A}/dC$ and also $\bar{\pi}_S$ from the Euler equation (9), the third line indicates how close we get to the true effect of climate from (8). As in Proposition 2, we see three sources of bias: the inability to identify effects of K , the Ω introduced by preparatory actions, and the $D/[h'(\bar{A})\bar{\chi}]$ that captures historical restraints and is the analogue of ω . The data will still not allow us to address the first, but we will eliminate the second and sign the third.

Consider the bias Ω . The following corollary shows that we can use the coefficient $\hat{\gamma}_0$ on longer-horizon forecasts to estimate Ω :

Corollary 6. *Let the conditions of Proposition 5 hold. Then:*

$$\frac{D}{h'(\bar{A})\bar{\chi}} \bar{\pi}_A \frac{\Omega}{1-g} = \frac{\hat{\gamma}_0}{\hat{\lambda}_0} \hat{\Lambda}_1 \frac{\hat{\lambda}_0}{\hat{\lambda}_1} = \hat{\Lambda}_1 \frac{\hat{\gamma}_0}{\hat{\lambda}_1}.$$

Proof. See Appendix E.13. □

Longer-horizon forecasts matter only by inducing preparatory actions: their effects are modulated by Ψ in equation (11). We can therefore use their ratio with shorter-horizon forecasts to identify the bias from preparatory actions, adjusting for the ex-post adaptation term $\hat{\Lambda}_1 \hat{\lambda}_0 / \hat{\lambda}_1$ (see Proposition 5) that motivates the preparation in Ω . Using Proposition 5 and labeling pieces as in (16), we can then calculate:

$$\begin{aligned}
&\underbrace{\hat{\Lambda}_0 - \hat{\Lambda}_1 \frac{\hat{\lambda}_0}{\hat{\lambda}_1}}_{\text{direct effects}} - \underbrace{\frac{1-\beta}{\beta} \hat{\Lambda}_1 \frac{\hat{\lambda}_0}{\hat{\lambda}_1}}_{\text{ex-post adaptation}} \overbrace{\left[-\frac{1-\beta}{\beta} \hat{\lambda}_0 + \frac{1-\beta}{\beta} \hat{\Lambda}_1 \frac{\hat{\gamma}_0}{\hat{\lambda}_1} \right]}^{\text{ex-ante adaptation}} \\
&\quad \underbrace{\left[-\frac{1-\beta}{\beta} \hat{\lambda}_0 \right]}_{\text{estimated}} \underbrace{\left[+ \frac{1-\beta}{\beta} \hat{\Lambda}_1 \frac{\hat{\gamma}_0}{\hat{\lambda}_1} \right]}_{\Omega \text{ adjustment}} \\
&= \bar{\pi}_w + \frac{D}{h'(\bar{A})\bar{\chi}} \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g} \right] \frac{d\bar{A}}{dC} \Big|_{K \text{ fixed}}. \tag{18}
\end{aligned}$$

We have successfully eliminated the bias from Ω .

Now consider the multiplicative bias introduced by $D/[h'(\bar{A})\bar{\chi}]$, which is the analogue to ω from Proposition 2. This remaining bias affects the estimates of ex-post and ex-ante adaptation. The next corollary shows that we can sign it:

Corollary 7. *Let the conditions of Proposition 5 hold, so that $\Psi \propto \hat{\Lambda}_2/\hat{\Lambda}_1$. If $\Psi > 0$, then $D/[h'(\bar{A})\bar{\chi}] < 1$. If $\Psi = 0$, then $D/[h'(\bar{A})\bar{\chi}] = 1$. If $\Psi < 0$, then $D/[h'(\bar{A})\bar{\chi}] > 1$.*

Proof. See Appendix E.14 □

The sign of Ψ controls the bias from $D/[h'(\bar{A})\bar{\chi}]$, as it also did for ω . Proposition 5 showed that we can estimate the sign of Ψ from $\hat{\Lambda}_2/\hat{\Lambda}_1$.³⁵ We thus know whether $D/[h'(\bar{A})\bar{\chi}]$ dampens or inflates the adaptation channels. If $D/[h'(\bar{A})\bar{\chi}] < 1$, then adaptation to climate is greater than implied by responses to weather, as when adjustment costs constrain short-run responses more than long-run responses. In that case, the top line of (18) 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 $D/[h'(\bar{A})\bar{\chi}] > 1$, then 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. In that case, the top line of (18) and the estimated direct effects bound the effect of climate from either side. Either way, we have bounded the effect of climate if either K is fixed or K does not interact with A_t or S_t . And remarkably, we have done so without needing to observe either the stock or actions and without needing to assume particular functional forms for payoffs or stock accumulation.

The purely reduced-form approaches in Section 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 rather simple: Corollary 4 required either (i) agents to be myopic ($\beta = 0$) or (ii) actions to be independent over time ($g, \bar{\pi}_{AS} = 0$). Of course, the present section's calculations also exactly recover the effects of climate if these conditions are met (or if $\pi_S = 0$ as in Section 3), so the indirect least squares approach directly weakens the assumptions required by conventional reduced-form approaches without sacrificing anything in terms of identification.³⁶

³⁵This result formalizes a conjecture from Deschênes and Greenstone (2012) about how storage decisions should affect a distributed lag model: storage decisions are intertemporal substitutes, which can manifest as alternating signs in the estimated lags. (Although note, from (13), that if $g > 0$, then intertemporal substitutes can be consistent with positive Ψ and thus with a constant sign across lags' coefficients.)

³⁶Both calculations require that infrastructure either is fixed or does not interact with shorter-run adaptation decisions. Relaxing that constraint will require either data with variation in infrastructure or assumptions about how infrastructure interacts with other adaptation choices.

5.1 When Forecasts Are Not Observable

Forecasts will be readily observable in many applications with daily data, but they will be observable only in some applications with monthly or annual data (e.g., Shrader, 2020).³⁷ Therefore consider the following, generically feasible regression, which uses leads of weather as proxies for forecasts:

$$\pi_{jt} = \alpha_j + \sum_{i=-2}^2 \Phi_i w_{j(t-i)} + \eta_{jt}. \quad (19)$$

The right-hand side contains only the fixed effect, the contemporary effect of weather, two lags of weather, and two leads of weather. Reprising the same steps we just followed, Appendix D shows that if the conditions of Proposition 5 hold and Σ is diagonal,³⁸ then

$$\begin{aligned} & \overbrace{\hat{\Phi}_0 - \frac{\hat{\Phi}_1}{\frac{\hat{\Phi}_2}{\hat{\Phi}_1} - \frac{1}{\beta}} + \frac{1}{\beta} \hat{\Phi}_{-1} + \frac{1}{\beta^2} \hat{\Phi}_{-2}}^{\text{direct effects}} \\ & \overbrace{-\frac{1-\beta}{\beta} \left[\frac{\hat{\Phi}_1}{\frac{\hat{\Phi}_2}{\hat{\Phi}_1} - \frac{1}{\beta}} - \hat{\Phi}_{-1} \frac{\hat{\Phi}_2}{\hat{\Phi}_1} - \frac{1}{\beta} \hat{\Phi}_{-2} \frac{\hat{\Phi}_2}{\hat{\Phi}_1} \right]}^{\text{ex-post adaptation}} \\ & \overbrace{-\frac{1-\beta}{\beta} \left[\hat{\Phi}_{-1} - \left(\frac{\hat{\Phi}_2}{\hat{\Phi}_1} - \frac{1}{\beta} \right) \hat{\Phi}_{-2} \right]}^{\text{ex-ante adaptation (estimated)}} \frac{1}{\Sigma_{22}/\text{trace}(\Sigma)} \\ & \overbrace{+\frac{1-\beta}{\beta} \frac{\hat{\Phi}_{-2}}{\hat{\Phi}_{-1}} \left[\frac{\hat{\Phi}_1}{\frac{\hat{\Phi}_2}{\hat{\Phi}_1} - \frac{1}{\beta}} - \hat{\Phi}_{-1} \frac{\hat{\Phi}_2}{\hat{\Phi}_1} - \frac{1}{\beta} \hat{\Phi}_{-2} \frac{\hat{\Phi}_2}{\hat{\Phi}_1} \right]}^{\text{ex-ante adaptation } (\Omega \text{ adjustment})} \frac{\Sigma_{22}/\text{trace}(\Sigma)}{\Sigma_{33}/\text{trace}(\Sigma)} \\ & = \bar{\pi}_w + \frac{D}{h'(\bar{A})\bar{\chi}} \left[\bar{\pi}_A + \bar{\pi}_S \frac{h'(\bar{A})}{1-g} \right] \frac{d\bar{A}}{dC} \Big|_{K \text{ fixed}}. \quad (20) \end{aligned}$$

The intuition for identification is largely the same as described following Proposition 5 and Corollary 7. The adjustments are here more complicated because forecasts now act as omitted variables that affect weather variables' coefficients (see Appendix D). Whereas before the ratio of forecasts was critical to eliminating Ω from the ex-ante adaptation channel, the ratio of leads now plays that role. Comparing to the true

³⁷Of course, this problem is not unique to the indirect least squares approach, as the regressions in Section 4.3 also used forecasts.

³⁸ Σ diagonal implies that weather shocks are serially uncorrelated within a location. We will see that this is a good approximation in the empirical application.

effect of climate in (8), we are again left with biases from $D/[h'(\bar{A})\bar{\chi}]$ and long-lived infrastructure, exactly as in (18). The bias from $D/[h'(\bar{A})\bar{\chi}]$ again inflates or deflates the estimates of ex-post and ex-ante adaptation. Analogously to prior results, Appendix D shows that $\hat{\Phi}_2/\hat{\Phi}_1$ is proportional to Ψ . From Corollary 7, we use this sign to learn about the direction of bias from $D/[h'(\bar{A})\bar{\chi}]$: the bias dampens the adaptation channels if $\hat{\Phi}_2/\hat{\Phi}_1 > 0$ and inflates them otherwise.

The most substantive difference between (20) and (18) is that our estimate of ex-ante adaptation from the coefficient $\hat{\Phi}_{-1}$ on the lead of weather tends to be too small in magnitude: $\hat{\Phi}_{-1}$ reflects the total variation in weather, but only a fraction $\Sigma_{22}/\text{trace}(\Sigma)$ of that variation was forecasted one period ahead of time. The bias from proxying forecasts with the lead of weather vanishes as the fraction goes to 1 because forecasts are then perfect. In contrast, if time $t+1$ weather is largely unknown at time t , then we may estimate very little ex-ante adaptation even though an agent would undertake substantial ex-ante adaptation to climate change. Analogous bias arises in our correction for Ω , which depends on the first and second leads of weather. The primary cost of proxying forecasts by the leads of weather is having to calibrate Σ to outside data so that we can mechanically undo these biases. Importantly, needing Σ to be diagonal and calibrated to outside data is a far weaker assumption than required for any of the reduced-form approaches in Section 4.3 to successfully recover the effects of climate—and we do not even need assumptions about Σ in environments with observable forecasts.

6 Estimating Climate Impacts in U.S. Agriculture

I now demonstrate the applicability of the new approach by re-analyzing a seminal application in the weather-climate literature, the impacts of climate change on U.S. agricultural profits (Deschênes and Greenstone, 2007).

The construction of the data follows an updated version of the methodology in Deschênes and Greenstone (2007) and Fisher et al. (2012). I have observations of county-level agricultural profits and acreage every 5 years from 1987 through 2017 from the U.S. Census of Agriculture. I follow previous literature in studying a measure of growing season degree days (i.e., accumulated heat within a temperature range favorable to plant growth), a measure of extreme growing season degree days (i.e., accumulated extreme heat, generally harmful to plant growth), and growing season precipitation. The preferred specification includes USDA Farm Resource Region-by-year fixed effects (as in Deschênes and Greenstone, 2012),³⁹ weights counties by average acreage (as in Deschênes and Greenstone, 2007), clusters standard errors by state (as in Fisher et al., 2012), and restricts the sample to counties east of the 100th

³⁹Appendix B provides further details and reports the variance explained by the weather variables (see Fisher et al., 2012). Appendix F.1 assesses sensitivity to instead defining regions as individual states (as in Deschênes and Greenstone, 2007) or as the whole country (as in Fisher et al., 2012).

meridian, which are less likely to be irrigated (Schlenker et al., 2005; Fisher et al., 2012). Appendix B further details the data, and Appendix F contains a variety of robustness checks.

I compare two different calculations of climate impacts. The *reduced-form approach* requires strong assumptions on the decision-making environment, whereas the *theory-based approach* bounds the effect of climate under far weaker conditions. First, following the spirit of previous literature and regression (7), I estimate

$$\pi_{ct} = \alpha_c + \psi_{rt} + \sum_{k=1}^K \theta^k w_{ct}^k + \eta_{ct}, \quad (21)$$

where c indicates counties, t indicates years, π_{ct} is agricultural profits, the α_c are county fixed effects, the ψ_{rt} are region-year fixed effects, and superscript k indexes weather variables of interest. The reduced-form approach's calculation of climate change's consequences multiplies each θ^k by the projected change in w^k over the course of the century. Proposition 1 showed that this approach recovers the effect of climate if there are no dynamic linkages (i.e., if $\pi_S = 0$). This approach therefore requires the absence of ex-ante adaptation and the independence of ex-post adaptation from any past decisions.

The theory-based approach implements the indirect least squares estimator of Section 5. This approach is valid even if $\pi_S \neq 0$. Following regression (19), I estimate

$$\pi_{ct} = \alpha_c + \psi_{rt} + \sum_{k=1}^K \sum_{i=-2}^2 \Phi_i^k w_{c(t-i)}^k + \eta_{ct}. \quad (22)$$

I use the resulting coefficients to calculate each term in (20) for each weather variable k . I then multiply the terms from (20) by the projected change in w^k over the course of the century. I also calculate $\hat{\Phi}_2^k / \hat{\Phi}_1^k$ in order to sign Ψ (using Proposition A-2 in Appendix D and Corollary 7) and thereby bound the effects of climate.

The theory-based calculations require a value for the discount factor β : I use an annual discount rate of 12%.⁴⁰ Further, because forecasts are unobserved in this application, I must follow Section 5.1 in assuming that weather is serially uncorrelated.⁴¹ The assumption of serially uncorrelated weather seems an acceptable starting point: over all U.S. counties from 1972 to 2019, the correlation between locally demeaned growing season degree days and its lag is 0.13, the correlation between locally demeaned extreme growing season degree days and its lag is 0.075, and the correlation between locally demeaned growing season precipitation and its lag is -0.014.

I calibrate Σ to the ability of global climate models to forecast average summer temperatures at the end of the previous summer. Becker et al. (2020) report

⁴⁰Appendix F.1 shows that results are not sensitive to the discount rate.

⁴¹These assumptions were unnecessary in the reduced-form approach because, following Section 3, it simply assumes away ex-ante adaptation and any other dynamic linkages by requiring $\pi_S = 0$. With that assumption, agents have no use for forecasts and discounting is irrelevant.

anomaly correlation coefficients that correspond to $[(\Sigma_{22} + \Sigma_{33})/\text{trace}(\Sigma)]^{1/2}$.⁴² Emily Becker kindly reanalyzed their data to weight locations by the farmland acreage variable used here. Depending on the precise target months and lead time chosen, the anomaly correlation coefficient over 1991–2020 is between 0.25 and 0.45 in a six-model ensemble. Using 0.35 and assuming that 5/6 of the reported correlation reflects information available only one year in advance, we have $\Sigma_{22}/\text{trace}(\Sigma) = 0.0851$ and $\Sigma_{33}/\text{trace}(\Sigma) = 0.0034$. I assess robustness to a range of plausible values in Appendix C.

6.1 Effects of Marginally Increased Average Weather

The top panel of Table 1 reports the reduced-form coefficients from regressions (21) and (22). Profits increase in same-year growing degree days (“GDD”), but both same-year and previous-year extreme growing degree days reduce profits (“Extreme GDD”). The central estimate suggests that same-year precipitation reduces profits, but this effect could easily go the other way. The signs of the central estimates alternate from the first to the second lag for both extreme growing degree days and precipitation. Most leads of weather do not have statistically significant effects on profits.

The lower panel of Table 1 reports the medians and, in parentheses, lower and upper quartiles for the combinations of model primitives from equation (20).⁴³ The signs of the direct effects and the ex-post adaptation effects are nearly all consistent with the signs of same-year and previous-year impacts on profits. The exception is that the sign of the direct effects of non-extreme growing degree days opposes the sign of same-year impacts on profits—I discuss this in more detail in Section 6.3. The ex-ante adaptation effects are noisily estimated and often not clearly different from zero, as were the coefficients on the leads of weather. Finally, the interquartile range for ex-post adaptation is negative for all three weather variables. This result suggests that $\pi_S \neq 0$ and thus that the reduced-form approach’s calculations will not successfully recover climate impacts.

6.2 Long-Run vs Short-Run Adaptation

The final row of Table 1 reports the $\hat{\Phi}_2^k/\hat{\Phi}_1^k$, whose sign matches the sign of Ψ . Inequality (5), which guarantees convergence of expected actions, implies $|\hat{\Phi}_2^k/\hat{\Phi}_1^k| < 1$.⁴⁴ Reassuringly, the estimates in Table 1 are consistent with the magnitude of $\hat{\Phi}_2^k/\hat{\Phi}_1^k$

⁴²The correlation is $\text{Cov}[w_t, f_{1,t-1}]/(\text{Var}[w_t]\text{Var}[f_{1,t-1}])^{1/2}$, where $\text{Cov}[w_t, f_{1,t-1}] = \text{Var}[f_{1,t-1}] = \Sigma_{22} + \Sigma_{33}$ and $\text{Var}[w_t] = \text{trace}(\Sigma)$.

⁴³I obtain these statistics by sampling from distributions defined by the means and covariance matrix from regression (22). The lower panel does not report means and standard errors because the distributions can be skewed due to division by coefficients.

⁴⁴Lemma 2 in Appendix E.6 shows that (5) implies $|\bar{Z}h'(\bar{A}) + g| < 1$, and the proof of Proposition A-2 in Appendix E.16 shows that $\bar{Z}h'(\bar{A}) + g = \hat{\Phi}_2/\hat{\Phi}_1$.

Table 1: Top: Estimated coefficients and standard errors from regressions (21) and (22). Bottom panel: Model primitives estimated by combining regression (22) with equation (20), reported as the median and lower/upper quartiles.

	GDD	Extreme GDD	Precip
<i>Reduced-Form Coefficients</i>			
$\hat{\theta}$	15 (8.6)	-120 (49)	-5.2 (3.2)
$\hat{\Phi}_0$	11 (8.7)	-91 (50)	-3.2 (3.2)
$\hat{\Phi}_1$	-8.2 (6.9)	-48 (25)	-6.8 (3.1)
$\hat{\Phi}_2$	-8.4 (10)	17 (26)	0.78 (3.1)
$\hat{\Phi}_{-1}$	2.2 (4.5)	24 (23)	3.2 (2.1)
$\hat{\Phi}_{-2}$	-12 (5.7)	-34 (22)	-1.1 (2.1)
<i>Theory-Implied Effects From (20)</i>			
Direct Effects	-4.7 (-19,6.2)	-140 (-180,-110)	-6.9 (-10,-4.0)
Ex-Post Adaptation	-1.9 (-4.8,-0.22)	-3.5 (-5.4,-1.5)	-0.83 (-1.1,-0.59)
Ex-Ante Adaptation (Estimated)	0.23 (-1.1,1.0)	3.2 (0.45,6.4)	-0.24 (-0.51,0.0092)
Ex-Ante Adaptation (Ω Adjustment)	-1.7 (-15,7.5)	-3.3 (-6.8,-0.54)	-0.20 (-0.66,0.12)
$\hat{\Phi}_2/\hat{\Phi}_1$	0.67 (-0.18,1.8)	-0.36 (-0.67,-0.0073)	-0.11 (-0.44,0.21)

All specifications include county and Farm Region-year fixed effects. The reduced-form estimates' standard errors and the theory-implied results derive from covariance matrices that are robust to clustering at the state level. The sample includes only counties east of the 100th meridian. Observations are weighted by (the square root of) a county's average farmland acreage. There are 16254 county-year observations and 37 state observations. Profits in thous. year 2002 dollars, GDD in °C-days, and precip in mm.

being less than 1. We have a case of intertemporal substitutes (complements) if this ratio is negative (positive). The estimates for non-extreme growing degree days and precipitation have ambiguous sign, but even the 75th percentile estimate is negative for the extreme growing degree days thought to drive climate impacts. Finding $\hat{\Phi}_2^k/\hat{\Phi}_1^k < 0$ is contrary to Le Châtelier's principle but consistent with recent empirical work in agricultural economics (Hendricks et al., 2014; Kim and Moschini, 2018). Following Eckstein (1984), these researchers attribute their results to soil nitrogen and pest dynamics inducing farmers to rotate their crops over time. Within the present paper's model, finding $\hat{\Phi}_2^k > 0$ with $\hat{\Phi}_1^k < 0$ implies that adaptive actions taken two years ago increase current payoffs by constraining the actions taken last year.

Table 2 explores the robustness of the estimated $\hat{\Phi}_2^k/\hat{\Phi}_1^k$.⁴⁵ The first row repeats the results from the preferred specification. The second row does not weight observations by farm acreage, the third and fourth rows explore alternate region-year fixed effects, the fifth row uses only years since 1997 in order to avoid an issue with older data (described in Appendix B), and the sixth row changes the sample to counties west of the 100th meridian. In most of these case, even the 75th percentile for extreme growing degree days is negative. The exceptions are cases with especially noisy estimates, whether because of limited identifying variation in the presence of state-year fixed effects or because of a restricted sample; however, the median estimates are still negative in even these cases. The seventh row estimates a geometric lag structure, using three lags and a one-step GMM estimator. The geometric term is equal to $\hat{\Phi}_2^k/\hat{\Phi}_1^k$. The estimate for extreme growing degree days is largely unchanged from the preferred specification. On the whole, the evidence supports $\hat{\Phi}_2^k/\hat{\Phi}_1^k < 0$ for extreme growing degree days.

The final two rows of Table 2 change the dependent variable from profits to yields. The theoretical analysis is for a maximand such as profits, but some argue that agents roughly act to maximize yields for given crop acreage. An advantage of using yields is the far greater number of observations available, as data are published annually instead of quinquennially. Corn yields' analogue of $\hat{\Phi}_2^k/\hat{\Phi}_1^k$ is negative. Soybean yields are the one case where we see a positive median estimate for extreme growing degree days. Because soybeans replenish soil nitrogen whereas corn depletes it, finding a positive estimate for soybean yields and a negative estimate for corn yields is consistent with soil nitrogen dynamics driving crop rotation.

6.3 Effects of Climate Change

I have thus far considered the marginal effects of non-extreme growing degree days, extreme growing degree days, and precipitation. I now multiply these effects by the projected changes due to climate change over the century in order to obtain a first-order approximation to the effects of climate change on the average acre of farmland. I project the effects of climate change using the RCP 4.5 trajectory of stabilized

⁴⁵The table does not vary the discount factor because doing so does not affect $\hat{\Phi}_2^k/\hat{\Phi}_1^k$.

Table 2: Robustness of $\hat{\Phi}_2^k/\hat{\Phi}_1^k$. Except where indicated, all specifications are as in the notes on Table 1.

	GDD	Extreme GDD	Precip
Base	0.67 (-0.18,1.8)	-0.36 (-0.67,-0.0073)	-0.11 (-0.44,0.21)
No Weighting	1.3 (0.4,2.9)	-0.87 (-1.8,-0.28)	-0.40 (-0.7,-0.14)
Year f.e.	0.77 (-3.0,3.2)	-0.56 (-0.73,-0.41)	-0.25 (-0.47,-0.055)
State-Year f.e.	-0.29 (-0.89,0.32)	-0.0023 (-2.1,2.4)	0.17 (-1.1,1.4)
1997–2017 Only	0.18 (-0.17,0.54)	-0.66 (-1.7,0.34)	-0.59 (-1.6,0.34)
Western U.S.	-0.43 (-0.64,-0.036)	-2.6 (-5.1,-0.76)	2.0 (1.2,3.6)
Three Lags ^a	0.65 (0.21,1.1)	-0.32 (-0.55,-0.095)	-1.2 (-1.4,-0.94)
Corn Yields ^b	-0.49 (-0.98,-0.056)	-1.4 (-2.8,0.028)	1.7 (1.3,2.3)
Soybean Yields ^b	0.29 (0.16,0.43)	1.1 (0.12,2.5)	-0.61 (-0.97,-0.35)

^a Interquartile range calculated from standard error.

^b Using annual data from 1987–2017.

Table 3: The percentage change in eastern U.S. agricultural profits due to predicted end-of-century changes in growing degree days, extreme growing degree days, and precipitation. The reduced-form estimates report central estimate and standard error. The theory-implied estimates report median and lower/upper quartiles.

	GDD	Extreme GDD	Precip	Combined
<i>Reduced-Form</i>				
	37	-79	-0.54	-42
	(21)	(33)	(0.33)	(22)
<i>Theory-Implied</i>				
Direct Effects	-12	-98	-0.71	-113
	(-46,15)	(-123,-72)	(-1.0,-0.41)	(-145,-82)
Ex-Post Adaptation	-4.6	-2.4	-0.086	-7.2
	(-12,-0.54)	(-3.7,-1.0)	(-0.11,-0.061)	(-15,-1.5)
Ex-Ante Adaptation	-90	-40	-0.91	-138
	(-899,470)	(-95,23)	(-1.8,-0.29)	(-1067,602)
Combined Adaptation	-92	-42	-0.99	-142
	(-905,466)	(-98,22)	(-1.9,-0.36)	(-1076,595)

All specifications include county and Farm Region-year fixed effects. The reduced-form estimates' standard errors and the theory-implied results derive from covariance matrices that are robust to clustering at the state level. The sample includes only counties east of the 100th meridian. Observations are weighted by (the square root of) a county's average farmland acreage. Climate projections use the RCP 4.5 scenario averaged across 21 CMIP5 models. There are 16254 county-year observations and 37 state observations.

emissions from 21 downscaled CMIP5 models. In this scenario, global mean surface temperature increases by around 2 degrees Celsius over the century, which increases growing degree days of both types (see Appendix B).

The top panel of Table 3 reports the conventional reduced-form calculation, which Proposition 1 showed was valid if $\pi_S = 0$. The projected increase in non-extreme growing degree days is estimated to increase agricultural profits, but the projected increase in extreme growing degree days reduces profits to a greater degree. Climate change reduces profits from the average acre of farmland by 42% in the central estimate, and widespread Le Châtelier intuition would suggest that this is an upper bound on the cost of climate change because adaptation will be greater in the long run. However, Table 1 showed that the coefficient on lagged extreme growing degree days appears to be nonzero, which suggests $\pi_S \neq 0$ and a need for the theory-based calculations.

The lower panel reports the new, theory-based estimates of climate impacts. The median direct effect projects losses of around 100% from climate change, which is over twice the estimate from the reduced-form approach. Effects on extreme growing degree days again drive the total effect of climate change. The primary source of the

difference with respect to the reduced-form approach is that additional non-extreme growing degrees here have harmful direct effects. As discussed around Proposition 2, the coefficient on contemporary weather in the reduced-form regression entwines the direct effect of weather with the immediate payoffs from ex-post adaptation to contemporary weather. The indirect least squares estimator cleans the coefficient $\hat{\lambda}_0$ of the immediate payoffs from ex-post adaptation. The remainder indicates that the direct effects are harmful.

Table 3 also shows that ex-post adaptation increases costs from climate change. How can adaptation reduce payoffs? In (8), an increase in \bar{A} increases payoffs from climate change if and only if $\bar{\pi}_A + \bar{\pi}_S h'(\bar{A})/(1-g) > 0$. The first term captures the immediate payoffs from adaptation and the second term captures the dynamic effects. The Euler equation (9) implies

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

The effects of adaptation on long-run payoffs from climate change therefore run counter to the effects on immediate payoffs defined by $\bar{\pi}_A$. Consider an example with $\bar{\pi}_A > 0$ and $h'(\bar{A}) > 0$. A dynamically optimizing agent would forgo immediate gains from further increasing A_t if simultaneously increasing the stock variable imposes costs in later periods. Moreover, because she is impatient, those later periods' non-discounted costs must exceed the immediate benefits that she forsakes around an optimum. When we examine the effects of climate change on steady state payoffs, we account for changes in both \bar{A} and \bar{S} without regard to this timing. The future costs of the larger stock dominate the calculation. In sum, detecting long-run costs from adaptation indicates that agents undertook actions that provided short-run benefits but left them a less desirable stock for the long run.

Now consider whether agents adapt to climate as they do to short-run weather shocks. The bias from $D/[h'(\bar{A})\bar{\chi}] \neq 1$ in (20) reflects historical restraints (as with ω in Section 4.3). If $\hat{\Phi}_2^k/\hat{\Phi}_1^k > 0$, the present calculations underestimate adaptation to climate (because $\Psi > 0$ and $D/[h'(\bar{A})\bar{\chi}] < 1$), but if $\hat{\Phi}_2^k/\hat{\Phi}_1^k < 0$, the present calculations overestimate adaptation to climate (because $\Psi < 0$ and $D/[h'(\bar{A})\bar{\chi}] > 1$). We saw in Table 1 that $\hat{\Phi}_2^k/\hat{\Phi}_1^k < 0$ for the extreme growing degree days that drive climate impacts. This result implies that we observe more adaptation to short-run weather shocks than would occur in response to long-run changes in climate. The implied resource scarcity story is intuitively consistent with finding that adaptation provides short-run benefits but imposes long-run costs.⁴⁶ Because $\hat{\Phi}_2^k/\hat{\Phi}_1^k < 0$, we can bound the effects of climate by the estimated total effects that include projected adaptation and by the estimated direct effects that exclude adaptation.

⁴⁶Aragón et al. (2021) show that Peruvian farmers increase acres planted in response to hot weather shocks. They speculate that these decisions will reduce future land productivity through the types of soil dynamics described in Section 6.2 as implying $\hat{\Phi}_2^k/\hat{\Phi}_1^k < 0$.

However, ex-ante adaptation is here imprecisely estimated. The row for ex-ante adaptation in Table 3 includes the adjustments for $\Sigma_{22}/\text{trace}(\Sigma)$ and $\Sigma_{33}/\text{trace}(\Sigma)$. The median estimates suggest that ex-ante adaptation is costly, but the estimates are noisy (especially for non-extreme growing degree days). Appendix C shows that the combination of ex-post and ex-ante adaptation appears to be costly as long as Σ_{33} is much smaller than Σ_{22} , which is a reasonable calibration. In this case, the direct effects of climate change are a lower bound on the total costs of climate change. This lower bound implies the complete or near-complete elimination of profits from the average acre of farmland. Further, this lower bound is only affected by the potential for changes in long-lived infrastructure to interact with shorter-run adaptation decisions if these interactions not only oppose the estimated effect of adaptation but do so strongly enough to flip the sign of climate's effect on adaptation actions in (10).⁴⁷ As there is no reason to believe such an extreme outcome is likely, the estimated direct effects partially identify the effects of climate from panel variation in weather.

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' decisions. I have shown that we can bound long-run effects by using a new indirect least squares estimator, and I have shown that the new estimator can generate very different conclusions than 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 a tight two-sided bound.

Instead of writing down a model of everything, I have highlighted the dynamic differences between transient weather shocks and permanent shifts in climate. Of course, weather shocks and climate change differ in other ways, including 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. In addition, I have followed the empirical literature in estimating the effects of changing one stationary climate to another. Future work should consider the process of changing the climate. By imposing stronger assumptions on the decision-making environment and constraining its parameters to replicate the long-run costs implied by the methods presented here, future work could simulate counterfactual climate trajectories and estimate the costs of transitioning from one climate to another.

⁴⁷And even then the most likely case is that direct effects provide a fairly tight upper bound on costs.

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 (2018a) “Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption using big data,” Working Paper 24397, National Bureau of Economic Research.
- (2018b) “Quantifying economic damages from climate change,” *Journal of Economic Perspectives*, Vol. 32, No. 4, pp. 33–52.
- Becker, Emily, Ben P. Kirtman, and Kathy Pegion (2020) “Evolution of the North American Multi-Model Ensemble,” *Geophysical Research Letters*, Vol. 47, No. 9, p. e2020GL087408.
- Blakeslee, David, Ram Fishman, and Veena Srinivasan (2020) “Way down in the hole: Adaptation to long-term water loss in rural India,” *American Economic Review*, Vol. 110, No. 1, pp. 200–224.
- 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 A, Amir Jina, Michael T Delgado, Michael Greenstone, Trevor Houser, Solomon M Hsiang, Andrew Hultgren, Robert E Kopp, Kelly E McCusker, Ishan B Nath, James Rising, Ashwin Rode, Hee Kwon Seo, Arvid Viaene, Jiacan Yuan, and Alice Tianbo Zhang (2020) “Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits,” Working Paper 27599, National Bureau of Economic Research, Series: Working Paper Series.
- 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.

- 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.
- Eckstein, Zvi (1984) “A rational expectations model of agricultural supply,” *Journal of Political Economy*, Vol. 92, No. 1, pp. 1–19.
- 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.
- 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.
- Hagerty, Nick (2020) “Adaptation to water scarcity in irrigated agriculture,” working Paper.
- Hamermesh, Daniel S. (1995) “Comment,” *ILR Review*, Vol. 48, No. 4, pp. 835–838.

- 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.
- Hendricks, Nathan P., Aaron Smith, and Daniel A. Sumner (2014) “Crop supply dynamics and the illusion of partial adjustment,” *American Journal of Agricultural Economics*, Vol. 96, No. 5, pp. 1469–1491.
- 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.
- Hornbeck, Richard and Pinar Keskin (2014) “The historically evolving impact of the Ogallala aquifer: Agricultural adaptation to groundwater and drought,” *American Economic Journal: Applied Economics*, Vol. 6, No. 1, pp. 190–219.
- 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 II: Empirical results,” *The Journal of Human Resources*, Vol. 37, No. 3, pp. 600–622.
- 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.
- Kim, Hyunseok and GianCarlo Moschini (2018) “The dynamics of supply: U.S. corn and soybeans in the biofuel era,” *Land Economics*, Vol. 94, No. 4, pp. 593–613.

- 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.
- Lazo, Jeffrey K., Rebecca E. Morss, and Julie L. Demuth (2009) “300 billion served: Sources, perceptions, uses, and values of weather forecasts,” *Bulletin of the American Meteorological Society*, Vol. 90, No. 6, pp. 785–798.
- 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.
- 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.
- Masseti, 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.
- Neidell, Matthew (2009) “Information, avoidance behavior, and health: The effect of ozone on asthma hospitalizations,” *Journal of Human Resources*, Vol. 44, No. 2, pp. 450–478.
- 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.
- 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.
- Sorkin, Isaac (2015) “Are there long-run effects of the minimum wage?” *Review of Economic Dynamics*, Vol. 18, No. 2, pp. 306–333.
- 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 für Nationalökonomie*, vol. 1, 1930, pp. 669–679),” 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.