

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 December 2019

I thank seminar participants at the 2019 Spring Meeting of the National Bureau of Economic Research Energy and Environment Program, the Bank of Canada Climate Change and Central Banking Workshop, the Queen's Institute of Intergovernmental Relations State of the Federation Workshop, Arizona State University, Auburn University, University of Alabama, University of Arizona, University of California Berkeley, University of Chicago, University of Massachusetts, University of Southern California, and University of Wisconsin for helpful comments. I thank Wolfram Schlenker for kindly providing data and Olivier Deschênes 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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
JEL No. D84,H43,Q54

ABSTRACT

I formally relate the consequences of climate change to the time series variation in weather extensively explored by recent empirical literature. I show that reduced-form fixed effects estimators can recover the effects of climate if agents are myopic, if agents' payoff functions belong to a particular class, or if the actions agents take in each period do not depend on actions taken in previous periods. I also show how to recover structural estimates of climate change impacts from reduced-form weather regressions. Applying this new method, I find that an additional 2°C of global warming would reduce eastern U.S. agricultural profits by around 50% under the median estimates.

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 severely hampered economists' ability to evaluate policy. Recognizing that different locations have different climates, many economists have hoped to estimate the effects of climate change from the spatial correlation between climate and outcomes of interest (e.g., Mendelsohn et al., 1994; Schlenker et al., 2005; Nordhaus, 2006). However, any two locations differ along many dimensions, leading to concerns about omitted variables bias.¹ Fortunately, though, the same location does experience different weather at different times. Stimulated by Deschênes and Greenstone (2007), an explosively growing empirical literature estimates the consequences of a location happening to experience cooler-than-average or hotter-than-average weather.² These researchers project the consequences of climate change by combining their credibly estimated effects of weather with scientists' predictions about how climate change will alter the distribution of weather. Whether the estimated weather treatment is in fact a good proxy for the unobserved climate treatment has been the subject of much debate but little analysis.³

I here undertake the first formal analysis that precisely delineates what and how we can learn about climate impacts from weather impacts. A change in climate differs from a weather shock in being repeated period after period and in affecting expectations of weather far out into the future. Linking weather to climate therefore requires analyzing a dynamic model that captures the distinction between transient and permanent changes in weather. I study an agent (equivalently, firm) who is exposed to stochastic weather outcomes. The agent chooses actions (equivalently, investments) that suit the weather. Actions can be responses to realized weather ("ex-post adaptation") or can be proactive investments against future weather ("ex-ante adaptation"). The actions chosen in different periods may be complements or substitutes: when actions are intertemporal complements, choosing a high action in the previous period reduces the cost of choosing a high action today, but when actions are intertemporal substitutes, choosing a high action in the previous period increases the cost of choosing a high action today. The first case is consistent with adjustment costs, and the second case is consistent with actions that require scarce resources.⁴ When choosing actions,

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) discuss the strengths and weaknesses of relying on panel variation in weather. Few doubt that weather shocks are as-good-as-randomly assigned. For instance, Dell et al. (2014, 741) write that "the primary advantage of the new literature is identification", and Blanc and Schlenker (2017, 262) describe "weather anomalies" as "ideal right-hand side variables" because "they are random and exogenous".

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."

⁴Both types of stories exist in the literature. For instance, in studies of the agricultural impacts of climate

the agent knows the current weather, has access to forecasts of future weather, and relies on knowledge of the climate to generate forecasts of weather at longer horizons. A change in the climate alters the distribution of potential weather outcomes as well as the agent's expectation of future weather outcomes.

I derive the effects of climate change in terms of model primitives and express reduced-form fixed effects estimators in terms of these same model primitives. I show that reduced-form estimates of weather impacts can exactly recover the theory-implied effects of climate change on payoffs in a few special cases. Two are of particular interest. First, if actions are neither intertemporal substitutes nor intertemporal complements (so that current decisions are not directly affected by previous decisions), then empirical researchers can recover the effects of climate on actions by combining the estimated effects on actions of current weather, lagged weather, and forecasts. And because empirical researchers can recover the effects of climate on actions, they can also recover the effect of climate on payoffs. Second, researchers do not need to recover effects on actions if agents' payoff functions satisfy a particular condition. Empirical researchers can then recover the effects of climate from especially simple regressions that are consistent with standard practice to date.

However, many applications will not satisfy the special cases. I therefore also extend conventional regression frameworks to recover structural estimates of climate impacts. Because I formally derive the reduced-form regression coefficients in terms of model primitives, I can recover combinations of model primitives from these coefficients and then calculate the theory-implied effects of climate change.⁵ The identification is exactly the same as in the recent reduced-form literature (relying on within-unit variation in weather), but the regression specification and the use of the estimated coefficients both differ from standard practice.

I apply this new method to the seminal analysis of climate and agriculture from Deschênes and Greenstone (2007). The results suggest that the special cases do not hold in this application. I find that adaptation offsets some of the costs of extreme heat in the short run, but because adaptation imposes its own costs, adaptation adds to the costs of extreme heat in the long run. In contrast, adaptation to non-extreme heat imposes costs in the short run in exchange for adding to the long-run benefits of additional days with non-extreme

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.

⁵This approach is in the spirit of Marschak's Maxim. Heckman (2010, 359) writes, "Marschak's Maxim suggests that economists should solve well-posed economic problems with minimal assumptions. All that is required to conduct many policy analyses or to answer many well-posed economic questions are policy invariant combinations of the structural parameters that are often much easier to identify than the individual parameters themselves and that do not require knowledge of individual structural parameters." It is also related to sufficient statistics approaches (see Chetty, 2009) and to price theory (see Weyl, 2019). Throughout, I use "reduced-form" and "structural" in the way now common in empirical work rather than in accord with their original usage.

heat. Most adaptation is *ex post*, but there is evidence of *ex-ante* adaptation to extreme heat in more recent years. In total, the costs of additional days with extreme heat outweigh the benefits of more days with non-extreme heat: the median estimates suggest that the current century's warming (in the RCP 4.5 scenario of stabilized emissions) would reduce agricultural profits by 56% in the absence of adaptation and by 50% if agents adapt as they do to annual weather shocks.

But will agents display more or less adaptation to climate than to weather? The estimated regression coefficients imply that actions are intertemporal substitutes, as in resource scarcity stories. This finding is consistent with recent empirical results in agricultural economics (Hendricks et al., 2014; Kim and Moschini, 2018) and with implications of crop rotation dynamics (Eckstein, 1984) but is contrary to widespread intuition based on Le Châtelier's principle. According to the theoretical analysis, agents will undertake more adaptation to short-run weather shocks than to long-run climate change. I can therefore bound the effects of climate change by the no-adaptation and full-adaptation estimates. The median estimates suggest losses of 50–56% (\$20–22 billion annually, in year 2017 dollars) from warming over the century.

There has been remarkably little 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 recent empirical literature to climate change. The primary exception is an argument given in Hsiang (2016) and repeated in Deryugina and Hsiang (2017). The argument stipulates that a change in climate differs from a change in weather only by affecting beliefs about future weather. This difference in beliefs can matter for payoffs only if it affects an agent's chosen actions. However, the envelope theorem tells us that an optimizing agent's actions cannot have first-order consequences for payoffs. Therefore the effects of weather on payoffs exactly—and generically—identify the effects of climate on payoffs.

By formalizing the distinction between climate and weather in a dynamic environment, the present analysis highlights two weak points in this argument.⁶ First, it is true that a change in climate alters beliefs about future weather, but it is also true that a change in climate alters past weather and past actions. Past actions are predetermined variables from the perspective of an optimizing agent and thus do not drop out through the envelope theorem. Even myopic agents can respond differently to weather and climate. Second, in a dynamic model, the envelope theorem applies to the intertemporal value function, not to the per-period payoff function investigated by much empirical work. Optimized current actions can have first-order effects on current payoffs when those are offset by first-order effects on

⁶In an earlier expositional analysis, I showed how envelope theorem arguments can fail in a three-period model (Lemoine, 2017). The present work precisely analyzes the consequences of climate change in an infinite-horizon model, constructively shows which types of empirical estimates can be informative about the climate, and develops and applies a new approach to structurally estimating the consequences of climate change.

expected future payoffs. Only in a special class of payoff functions will optimized current actions not have first-order effects on future payoffs.

A few other lines of research are related. First, calibrated numerical simulations have shown that dynamic responses are critical to the effects of climate on timber markets (Sohn-gen and Mendelsohn, 1998; Guo and Costello, 2013) and to the cost of increased cyclone risk (Bakkensen and Barrage, 2018). I develop a general analytic setting that precisely disentangles several types of dynamic responses and relates them to widely used fixed effects estimators. Second, empirical work has shown that agents use forecasts of future weather, even at seasonal scales. In particular, Shrader (2017) and Taraz (2017) use variation in seasonal forecasts and in past years' weather outcomes, respectively, to identify ex-ante adaptation to weather events. I formally demonstrate that estimating responses to forecasts and lagged weather is critical to recovering the consequences of climate change. Finally, Kelly et al. (2005) and Kala (2017) study learning about the climate from observed weather. I here abstract from learning in order to focus on mechanisms more relevant to the recent empirical literature.

The challenge of attempting to estimate long-run effects from short-run variation is a common one in empirical economics. To get around this challenge, environmental economists have found policy-induced variation in long-run pollution exposure that is plausibly exogenous to health outcomes (e.g., Chen et al., 2013; Anderson, 2015; Barreca et al., 2017; Bishop et al., 2018). Unfortunately, this type of variation may not be available to researchers interested in the consequences of changing the climate. 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).⁷ And macroeconomists formerly hoped to learn about long-run output-inflation tradeoffs by estimating distributed lag models, but Lucas (1972) argued that, when agents have rational expectations, the lagged 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 derives the theory-implied effect of climate. Section 4 establishes conditions under which the effect of climate can be recovered from reduced-form estimates of weather impacts. Section 5 develops the new method of structurally estimating climate impacts and applies it to U.S. agriculture. The final section describes potential extensions. The appendix contains empirical details, proofs, and

⁷Three 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 (2019) show that estimated responses to transient shocks can differ substantially from the theory-implied causal effects that empirical researchers seek to test. The present paper is similar in deriving sufficient conditions for estimated effects to match theory-implied effects. Third, Keane and Wolpin (2002) describe tradeoffs between cross-sectional and time series variation when estimating the effects of welfare benefits. These tradeoffs are similar to those that motivate the present paper.

additional results, including formal analysis of recently popular long difference estimators.

2 Setting

An agent is repeatedly exposed to stochastic weather outcomes and takes actions based on realized weather and information about future weather. The realized weather in period t is w_t and the agent’s chosen action is A_t . This action may be interpreted as a level of activity (e.g., time spent outdoors, energy used for heating or cooling, irrigation applied to a field) or as a stock of capital (e.g., outdoor gear, size or efficiency of furnace, number or efficiency of irrigation lines). The agent’s time t payoffs are given by the twice-differentiable function $\pi(A_t, A_{t-1}, w_t, w_{t-1}) : \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$.⁸ Letting subscripts indicate partial derivatives, I assume $\pi_{11} < 0$ and $\pi_{22} \leq 0$, implying declining marginal benefits of current and past actions.

I interpret actions as adaptations that become more valuable with high weather outcomes ($\pi_{13}, \pi_{23} \geq 0$). Following terminology from the literature on climate adaptation (e.g., Fankhauser et al., 1999; Mendelsohn, 2000), a case with $\pi_{13} > 0$ reflects adaptation that can occur after weather is realized (“reactive” or “ex-post” adaptation) and a case with $\pi_{23} > 0$ reflects adaptation that can occur before weather is realized (“anticipatory” or “ex-ante” adaptation).⁹ I allow adaptation to play both roles at once. The possibility that $\pi_4 \neq 0$ reflects potential delayed impacts from the previous period’s weather, with π_{14} and π_{24} capturing the potential for ex-post adaptation to alter these delayed impacts. Consistent with the normalizations above, I assume $\pi_{14}, \pi_{24} \geq 0$. Finally, observe that the actions could reflect a firm’s production responses to price signals rather than responses to weather per se. In this interpretation, the normalizations imply that “high” weather outcomes increase the price of a firm’s output or reduce the cost of its input.

I allow π_{12} to be positive or negative, with its magnitude constrained as described below. When $\pi_{12} < 0$, actions are “intertemporal substitutes”, so that choosing a higher level of past actions increases the cost of choosing higher actions today. I describe this case as a resource scarcity story.¹⁰ For instance, pumping groundwater today raises the cost of pumping groundwater tomorrow (see Blanc and Schlenker, 2017) or rescheduling activities around today’s weather makes it hard to reschedule activities around tomorrow’s weather (see Graff Zivin and Neidell, 2009). When $\pi_{12} > 0$, actions are “intertemporal complements”,

⁸I generalize to vector-valued actions and multidimensional weather in the appendix. Doing so yields little new insight at the expense of exposition.

⁹When interpreting actions as the choice of capital stock, the payoff function is consistent with standard models of depreciation. If we restrict the payoff function to allow only ex-ante adaptation, then the setting corresponds to a time-to-build model with a one-period lag. Section 6 discusses the implications of capital stocks that take longer to build.

¹⁰Relating to the literature on resource extraction, the case with $\pi_{12} < 0$ can be seen as reflecting stock-dependent extraction costs (Heal, 1976).

so that choosing a higher level of past actions increases the benefit from choosing higher actions today. I describe this case as an adjustment cost story.¹¹ For instance, small changes to cropping practices or activity schedules may be easier to implement than large changes. The magnitude of π_{12} affects the agent's preferred timing of adaptation. As $|\pi_{12}|$ becomes large, the agent prefers to begin adapting before the weather event arrives, but when $|\pi_{12}|$ is small, the agent may wait to undertake most adaptation only once the weather event has arrived.¹²

The agent observes time t weather before selecting her time t action. The agent also understands 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 $t - 1$, the agent's only information about time t weather consists in knowledge of the climate. However, at time $t - 1$ the agent receives a forecast f_{t-1} of time t weather: $f_{t-1} = C + \zeta\nu_{t-1}$, where the innovation ν_{t-1} is a mean-zero, serially uncorrelated random variable with variance $\tau^2 > 0$. The forecast is an unbiased predictor of time t weather: $w_t = f_{t-1} + \zeta\epsilon_t$, where ϵ_t is a mean-zero, serially uncorrelated random variable with variance $\sigma^2 > 0$.¹³ The parameter $\zeta \geq 0$ is a perturbation parameter that will be useful for analysis (see Judd, 1996). The covariance between ϵ_t and ν_t is ρ . The covariance between w_t and w_{t-1} is then $\zeta^2\rho$. The agent incorporates knowledge of such serial correlation in her forecasts.

The agent maximizes the present value of payoffs over an infinite horizon:

$$\max_{\{A_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t E_0 [\pi(A_t, A_{t-1}, w_t, w_{t-1})],$$

where $\beta \in [0, 1)$ is the per-period discount factor, A_{-1} is given, and E_0 denotes expectations

¹¹The benchmark quadratic adjustment cost model has $\pi_{12} = k$ for some $k > 0$ (see Hamermesh and Pfann, 1996).

¹²The magnitude of π_{12} is related to the distinction between ex-post and ex-ante adaptation insofar as it affects the agent's preferred timing of adaptation actions. However, π_{12} incentivizes early adaptation only to reduce the costs of later adaptation, not because early adaptation provides protection from weather events. I reserve the terms ex-ante and ex-post adaptation to refer to the effects of actions on the marginal benefit of weather, captured by π_{13} , π_{23} , π_{14} , and π_{24} .

¹³Consistent 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 one period ahead is $\zeta^2(\sigma^2 + \tau^2)$, so we need to apportion any change in variance between σ^2 and τ^2 . I leave such an extension to future work.

at the time 0 information set. The solution satisfies the following Bellman equation:

$$V(A_{t-1}, w_t, f_t, w_{t-1}; \zeta) = \max_{A_t} \left\{ \pi(A_t, A_{t-1}, w_t, w_{t-1}) + \beta E_t [V(A_t, w_{t+1}, f_{t+1}, w_t; \zeta)] \right\}$$

$$\text{s.t. } w_{t+1} = f_t + \zeta \epsilon_{t+1}$$

$$f_{t+1} = C + \zeta \nu_{t+1}.$$

Time t payoffs do depend indirectly on actions and weather prior to time $t-1$ because earlier actions affect time $t-1$ actions.

The setting is sufficiently general to describe many applications of interest. For instance, much empirical literature has studied the effects of weather on energy use. The agent could then be choosing indoor temperature in each period, where payoffs depend on current actions through energy use and depend on weather through thermal comfort. Habituation to outdoor temperatures is captured by π_{14} . Much empirical work has also studied the effect of weather on labor productivity. The decision variable could be effort, the dependence of payoffs on weather could reflect current thermal stress as well as the effects of the previous day's weather via sleep and physiological functioning, the resource scarcity is one of tasks needing to be done, and forecasts allow the agent to plan tasks and vacation time around weather outcomes. Finally, many researchers have studied the effects of weather on agricultural outcomes. In that case, payoffs are profits, actions include planting decisions, and weather affects yields.

I will often impose one of the following two assumptions:

Assumption 1. ζ^2 is small.

Assumption 2. π is quadratic.

Either assumption 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).¹⁴

I will be interested in empirical researchers' ability to estimate the consequences of altering C from observable responses to time series variation in w_t and f_t . It is important to be clear about the climate experiment. I study the effects of a change in climate on an agent who has had time to adapt to the new climate. This climate change treatment is consistent with the exercise common in the empirical literature, which calculates the effect of replacing today's distribution of weather with a distribution projected to hold by the end of the century. I will not study how the transition from one climate to another interacts with agents' decisions,¹⁵ and I will not study how expectations of a future change in climate

¹⁴Note that when applying Assumption 2, the chosen policy is affected by the variance of weather (through the realized weather) even though the policy rule is independent of that variance.

¹⁵Kelly 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

affect agents today.¹⁶ 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 Theory-Implied Effect of Climate Change

I now derive the exact effect of climate change on long-run payoffs and actions within this model. I later explore how to estimate these effects from observable variation in weather.

The analysis approximates the solution to the full, stochastic model around the steady state of the deterministic model, which sets $\zeta = 0$ (Judd, 1996). The first-order condition for the deterministic model is:

$$0 = \pi_1(A_t, A_{t-1}, C, C) + \beta V_1(A_t, C, C, C; 0).$$

The envelope theorem yields:

$$V_1(A_{t-1}, C, C, C; 0) = \pi_2(A_t, A_{t-1}, C, C).$$

Advancing this forward by one timestep and substituting into the first-order condition, we have the Euler equation:

$$0 = \pi_1(A_t, A_{t-1}, C, C) + \beta \pi_2(A_{t+1}, A_t, C, C). \quad (1)$$

A steady state \bar{A} of the deterministic system is implicitly defined by

$$0 = \pi_1(\bar{A}, \bar{A}, C, C) + \beta \pi_2(\bar{A}, \bar{A}, C, C). \quad (2)$$

Define $\bar{\pi} \triangleq \pi(\bar{A}, \bar{A}, C, C)$. The following lemma describes the uniqueness and stability of the steady state.¹⁷

Lemma 1. *For $\bar{\pi}_{12} \neq 0$ and $\zeta = 0$, \bar{A} is locally saddle-path stable if and only if $(1 + \beta)|\bar{\pi}_{12}| < -\bar{\pi}_{11} - \beta\bar{\pi}_{22}$, in which case \bar{A} is unique. For $\bar{\pi}_{12} = 0$ and $\zeta = 0$, the agent chooses $A_t = \bar{A}$ at all times t .*

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.

¹⁶Severen 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 time series context and show how to use them to improve panel analyses in the tradition of Deschênes and Greenstone (2007).

¹⁷The steady state exists if $\sup_{\bar{A}}[\pi_1(\bar{A}, \bar{A}, C, C) + \beta\pi_2(\bar{A}, \bar{A}, C, C)] > 0$ and $\inf_{\bar{A}}[\pi_1(\bar{A}, \bar{A}, C, C) + \beta\pi_2(\bar{A}, \bar{A}, C, C)] < 0$. These conditions will hold in most applications.

Proof. See appendix. □

I henceforth assume that $(1 + \beta)|\bar{\pi}_{12}| < -\bar{\pi}_{11} - \beta\bar{\pi}_{22}$, so that the deterministic steady state is unique and saddle-path stable.

Now consider optimal actions in the stochastic system. The first-order condition is:

$$0 = \pi_1(A_t, A_{t-1}, w_t, w_{t-1}) + \beta E_t[V_1(A_t, w_{t+1}, f_{t+1}, w_t; \zeta)].$$

The envelope theorem yields:

$$V_1(A_{t-1}, w_t, f_t, w_{t-1}; \zeta) = \pi_2(A_t, A_{t-1}, w_t, w_{t-1}).$$

Advancing this forward by one timestep and substituting into the first-order condition, we have the stochastic Euler equation:

$$0 = \pi_1(A_t, A_{t-1}, w_t, w_{t-1}) + \beta E_t[\pi_2(A_{t+1}, A_t, w_{t+1}, w_t)]. \quad (3)$$

The following lemma describes the evolution of $E_0[A_t]$.

Lemma 2. *Let either Assumption 1 or 2 hold, and let $E_0[(A_1 - \bar{A})^2]$ be small. Then $\lim_{t \rightarrow \infty} E_0[A_t] = \bar{A}$.*

Proof. See appendix. □

When the conditions of Lemma 2 hold, applying the implicit function theorem to equation (2) yields:

$$\lim_{t \rightarrow \infty} \frac{dE_0[A_t]}{dC} = \frac{d\bar{A}}{dC} = \frac{\overbrace{\bar{\pi}_{13} + \bar{\pi}_{14} + \beta\bar{\pi}_{24}}^{\text{ex-post}} + \overbrace{\beta\bar{\pi}_{23}}^{\text{ex-ante}}}{-\bar{\pi}_{11} - (1 + \beta)\bar{\pi}_{12} - \beta\bar{\pi}_{22}} \geq 0. \quad (4)$$

This is the average long-run effect of climate change on actions. Expected future actions increase in the climate index because I normalize high actions to be more beneficial when the weather index is high. Equation (4) captures how climate change alters weather in all periods: the past, the present, and the future. We see the various forms of ex-post adaptation captured by $\bar{\pi}_{13}$, $\bar{\pi}_{14}$, and $\beta\bar{\pi}_{24}$. We also see the possibility of ex-ante adaptation, controlled by $\bar{\pi}_{23}$ and arising because the agent understands that the altered climate affects weather in subsequent periods. Finally, observe that $\bar{\pi}_{12}$ enters through the denominator in (4). When actions are intertemporal substitutes ($\bar{\pi}_{12} < 0$), this term reduces the magnitude of the response to climate change, as when resource scarcity makes long-run responses smaller than short-run responses. However, when actions are intertemporal complements ($\bar{\pi}_{12} > 0$), this term increases the magnitude of the response to climate change, as when adjustment costs allow long-run responses to exceed short-run responses.

Approximating the payoff function around the steady state, $w_t = w_{t-1} = C$, and $\zeta = 0$ and using either Assumption 1 or Assumption 2, we have:

$$\begin{aligned} E_0[\pi(A_t, A_{t-1}, w_t, w_{t-1})] &= \bar{\pi} + \bar{\pi}_1(E_0[A_t] - \bar{A}) + \bar{\pi}_2(E_0[A_{t-1}] - \bar{A}) \\ &\quad + \frac{1}{2}\bar{\pi}_{11}E_0[(A_t - \bar{A})^2] + \frac{1}{2}\bar{\pi}_{22}E_0[(A_{t-1} - \bar{A})^2] + \frac{1}{2}(\bar{\pi}_{33} + \bar{\pi}_{44})\zeta^2(\sigma^2 + \tau^2) \\ &\quad + \bar{\pi}_{12}E_0[(A_t - \bar{A})(A_{t-1} - \bar{A})] + \bar{\pi}_{13}Cov_0[A_t, w_t] + \bar{\pi}_{23}Cov_0[A_{t-1}, w_t] \\ &\quad + \bar{\pi}_{14}Cov_0[A_t, w_{t-1}] + \bar{\pi}_{24}Cov_0[A_{t-1}, w_{t-1}] + \bar{\pi}_{34}\zeta^2\rho, \end{aligned} \quad (5)$$

for $t > 1$. Differentiating equation (5) with respect to C and applying either Assumption 1 or Assumption 2 again, we find that

$$\lim_{t \rightarrow \infty} \frac{dE_0[\pi(A_t, A_{t-1}, w_t, w_{t-1})]}{dC} = \bar{\pi}_3 + \bar{\pi}_4 + [\bar{\pi}_1 + \bar{\pi}_2] \frac{d\bar{A}}{dC}. \quad (6)$$

The marginal effect of climate on long-run payoffs is composed of the direct effect of a larger weather index, in both the present ($\bar{\pi}_3$) and the past ($\bar{\pi}_4$), and the effects of changing long-run actions, including both present actions ($\bar{\pi}_1$) and past actions ($\bar{\pi}_2$). Equation (2) implies $\bar{\pi}_1 = -\beta\bar{\pi}_2$. Therefore,

$$\lim_{t \rightarrow \infty} \frac{dE_0[\pi(A_t, A_{t-1}, w_t, w_{t-1})]}{dC} = \bar{\pi}_3 + \bar{\pi}_4 + (1 - \beta)\bar{\pi}_2 \frac{d\bar{A}}{dC}. \quad (7)$$

Whether economic responses increase or decrease payoffs depends on the sign of $\bar{\pi}_2$. A case with $\bar{\pi}_2 > 0$ is a case in which higher actions impose costs today but provide benefits tomorrow, as when undertaking adaptation investments that take time to build. A case with $\bar{\pi}_2 < 0$ is a case in which higher actions provide benefits today but impose costs tomorrow, as when borrowing money or selling from storage. Undertaking more actions because of climate change increases payoffs if and only if actions are of the former type.

Equation (6) shows that changes in average payoffs depend on changes in actions. However, some previous literature has relied on envelope theorem intuition to argue that changes in actions are in fact irrelevant to effects on payoffs. Each solid curve in Figure 1 plots time t payoffs as a function of time t weather holding current actions, past actions, and past weather fixed. The three solid curves represent a continuum obtained by varying the agent's current actions. The dotted line gives the effect of time t weather on maximized time t payoffs, with points (a) and (b) indicating maximized payoffs at two particular weather realizations. The slope of the dotted line does not depend on how actions change with weather: because it is tangent to the solid lines, it matches the partial derivative of payoffs with respect to weather. This is the content of the envelope theorem. If a change in climate only affected current weather, then we could ignore how actions respond to the change in climate.

However, a change in climate also affects past and future weather. First consider the consequences of affecting past weather. Interpret weather as temperature. Imagine that

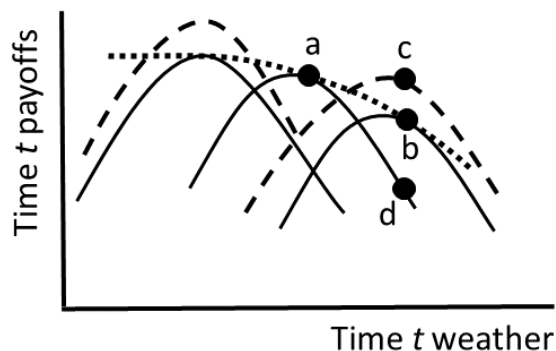


Figure 1: Per-period payoffs π_t against weather w_t . Each solid curve holds actions fixed. The dotted curve through points (a) and (b) gives the effect of time t weather shocks on myopically maximized payoffs. Point (c) shows the effect of weather shocks repeated in times $t - 1$ and t , and point (d) illustrates dynamically optimizing agents.

changing inputs imposes adjustment costs (i.e., $\pi_{12} > 0$) and that the solid curves are conditioned on previous actions having been chosen in response to experiencing the temperature at point (a). The dashed curves in Figure 1 plot payoffs conditional on past actions having been chosen in response to experiencing the indicated temperature, not the temperature at point (a). Payoffs increase because the preferred current actions do not impose as many adjustment costs when past actions are better-matched to current weather. Whereas a transient weather shock induces a change from point (a) to point (b), an enduring weather shock induces a change from point (a) to point (c). This figure illustrates intuition based on Le Châtelier's principle for why the effects of transient weather shocks on payoffs could overstate the effects of climate change (e.g., Deschênes and Greenstone, 2007). What does envelope theorem intuition miss? Changes in actions chosen at time $t - 1$ can have first-order effects on time t payoffs. Responses to transient weather shocks can fail to identify the effects of climate because they do not capture how climate affects the trajectory of previous actions.

Now consider the implications of climate affecting future weather. An agent who expects a transient weather shock to occur yet again in period $t + 1$ may choose actions that fail to maximize short-run payoffs if these actions increase payoffs in subsequent periods, as when $\pi_{23} \neq 0$. As an example, point (d) could be optimal even though its action maximizes short-run payoffs only at point (a). Further, if $\pi_{12} \neq 0$, agents may choose not to maximize short-run payoffs even when not expecting a time $t + 1$ weather shock (see equation (2)). What does envelope theorem intuition miss? Dynamic optimization typically requires trading off short-run and long-run effects. The envelope theorem no longer applies to per-period payoffs in the presence of such tradeoffs, instead applying only to expected present discounted payoffs.

4 Estimating the Effect of Climate Change from Reduced-Form Weather Regressions

I have derived the theory-implied long-run effect of climate change, but researchers do not know all of the structural parameters required to calculate this effect. Instead, empirical researchers have sought to estimate the effect of climate from reduced-form regressions with observed weather. I now consider whether and how such reduced-form regressions can recover the effect of climate.¹⁸

4.1 Estimating Effects on Actions

We have seen that the effects of climate on payoffs are closely related to its effects on actions. Further, much empirical research has sought to estimate the consequences of climate change for decision variables or functions of decision variables, including productivity (Heal and Park, 2013; Zhang et al., 2018), time allocation (Graff Zivin and Neidell, 2014), and energy use (Auffhammer and Aroonruengsawat, 2011; Deschênes and Greenstone, 2011; Auffhammer, 2018a). I therefore begin by considering the potential to estimate the effect of climate on actions from time series variation in weather.

First consider the determinants of time t actions. The proof of Lemma 2 shows that if either Assumption 1 or 2 holds and $(A_{t-1} - \bar{A})^2$ is small, then

$$\begin{aligned}
 A_t = \bar{A} + & \underbrace{\frac{\bar{\pi}_{14}}{\chi_2}(w_{t-1} - C) + \frac{\bar{\pi}_{12}}{\chi_2}(A_{t-1} - \bar{A})}_{\text{effects of past weather}} + \underbrace{\frac{\bar{\pi}_{13} + \beta\bar{\pi}_{24} + \beta\bar{\pi}_{14}\frac{\bar{\pi}_{12}}{\chi_1}}{\chi_2}(w_t - C)}_{\text{effects of current weather}} \\
 & + \underbrace{\frac{\beta\bar{\pi}_{23} + \beta\left(\bar{\pi}_{13} + \beta\bar{\pi}_{24} + \beta\bar{\pi}_{14}\frac{\bar{\pi}_{12}}{\chi_0}\right)\frac{\bar{\pi}_{12}}{\chi_1}}{\chi_2}(f_t - C)}_{\text{effects of future weather}}, \tag{8}
 \end{aligned}$$

where each $\chi_i > |\pi_{12}|$. We see time t actions determined by past, present, and future weather. Figure 2 illustrates the main relationships identified by this expression.

Actions depend on present weather in three ways. First, actions respond to current weather as a means of mitigating its immediate harm or amplifying its immediate benefits. This channel is controlled by $\bar{\pi}_{13}$. Second, actions respond to current weather when current actions can mitigate the harm or amplify the benefits incurred by current weather in future periods. This channel is controlled by $\bar{\pi}_{24}$ and arises only for forward-looking agents. As an

¹⁸I here consider only the internal validity of estimated effects. Equations (4) and (6) imply that the effect of climate change will vary with the current climate unless weather enters π only linearly. Empirical researchers should therefore take care when extrapolating estimated effects across locations and when pooling data across locations.




Figure 2: The determinants of present actions, from equation (8).

example of the distinction between the two channels, an agent may avoid going outside on a cold day both to minimize discomfort from the current temperature and to avoid getting sick in the near future. Both of these channels are forms of ex-post adaptation. Third, when $\bar{\pi}_{14} \neq 0$, current weather will affect the agent's chosen action in the next period (not pictured in Figure 2), leading a forward-looking agent to adjust her current action in preparation for that choice. This channel vanishes when $\bar{\pi}_{12} = 0$ because today's actions then do not directly interact with subsequent actions.

Actions also depend on forecasts of future weather. When there is the possibility of ex-ante adaptation ($\bar{\pi}_{23} > 0$), the agent chooses today's actions in order to directly mitigate the consequences (or enhance the benefits) of expected future weather. Further, expected future weather also affects desired future actions. The agent takes preparatory actions today that help her achieve her desired future actions. When $\bar{\pi}_{12} > 0$, a high forecast leads the agent to choose high actions today as a means of reducing future adjustment costs, but when $\bar{\pi}_{12} < 0$, a high forecast leads the agent to choose low actions today as a means of conserving resources for the future.

Past weather also affects actions in two ways. First, past weather affects the marginal payoffs from current actions directly when $\bar{\pi}_{14} \neq 0$. This is a form of ex-post adaptation. Second, past weather affects past actions, which impose historical restraints on current actions when $\bar{\pi}_{12} \neq 0$. When actions are intertemporal complements ($\bar{\pi}_{12} > 0$), high past actions justify higher present actions as a way to reduce adjustment costs, but when actions are intertemporal substitutes ($\bar{\pi}_{12} < 0$), high past actions justify lower present actions by depleting the resources needed to maintain a high action.

Empirical researchers hope to recover (4) from time series variation in weather. Let there be J agents (equivalently, firms) observed in each of T periods. Index these agents by j . In order to focus on the issue at hand, imagine that they are in the same climate C with the same payoff function π and the same stochastic process driving forecasts and weather, though each agent draws its own sequence of weather and forecasts. Consider the following fixed effects regression:

$$A_{jt} = \alpha_j + \Gamma_1 w_{jt} + \Gamma_2 w_{j(t-1)} + \Gamma_3 f_{jt} + \Gamma_4 A_{j(t-1)} + \eta_{jt}, \quad (9)$$

where α_j is a fixed effect for unit j and η_{jt} is an error term that is uncorrelated with the

covariates.¹⁹ I use a hat to denote the probability limit of each estimator. The following proposition relates the estimated coefficients to the effect of climate change.

Proposition 1. *Let either Assumption 1 or 2 hold, and let $(A_{j(t-1)} - \bar{A})^2$ be small for all observations. Then $\hat{\Gamma}_4 \propto \bar{\pi}_{12}$ and*

$$\hat{\Gamma}_1 + \hat{\Gamma}_2 + \hat{\Gamma}_3 = \omega \left(\frac{d\bar{A}}{dC} + \beta \bar{\pi}_{12} \Omega \right), \quad (10)$$

where $\bar{\pi}_{12} > 0$ implies $\omega \in (0, 1)$, $\bar{\pi}_{12} < 0$ implies $\omega > 1$, $\bar{\pi}_{12} = 0$ implies $\omega = 1$, and $\Omega \propto \bar{\pi}_{13} + \bar{\pi}_{14} + \beta \bar{\pi}_{24} \geq 0$.

Proof. See appendix. □

The three coefficients capture the three temporal relationships altered by climate change: $\hat{\Gamma}_1$ recovers consequences of altering current weather, $\hat{\Gamma}_2$ recovers consequences of altering past weather, and $\hat{\Gamma}_3$ recovers consequences of altering expectations of future weather. However, we cannot in general recover the response to a permanent change in climate from the estimated response to transient weather shocks. The reason for this failure is the possibility that $\bar{\pi}_{12} \neq 0$, which occurs if and only if $\hat{\Gamma}_4 \neq 0$.

Relationships of intertemporal substitutability or complementarity drive two types of wedges between the estimator in (10) and the effect of climate change in (4). The second term in parentheses in (10) reflects preparatory actions that are undertaken in response to forecasts but are not relevant to the long-run effects of climate. The fixed effects estimator is identified from shocks to forecasts and weather. As described above, a high forecast increases present actions both through the possibility of ex-ante adaptation and through preparatory actions. The former are important components of the effect of climate but the latter are not: an increase in the climate index C does increase forecasts, but because it also increases current and past weather, preparatory actions are not relevant to its long-run effects. When $\bar{\pi}_{12} > 0$, preparatory actions make the fixed effects estimator overstate responses to climate as observed agents are motivated by expectations of temporary adjustment costs, but when $\bar{\pi}_{12} < 0$, preparatory actions make the fixed effects estimator understate responses to climate as observed agents temporarily conserve resources.

The second wedge in (10) arises from 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. When $\bar{\pi}_{12} > 0$, historical restraints prevent an agent from adjusting too much to a transient weather shock, but when

¹⁹I do not explicitly model the unobservable characteristics that motivate the fixed effects specification. I am here not interested in identification but in what we learn from a well-identified weather regression. See Dell et al. (2014) and Auffhammer (2018b), among others, for expositions of identification in the climate-economy literature. I assume that the only possible sources of omitted variables bias are the failure to control for variables such as forecasts and lagged actions that are defined within the theoretical model.

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. The $\omega < 1$ captures how responses to transient shocks overstate historical restraints in this case. Consistent with conjectures in Deschênes and Greenstone (2007), observable short-run responses are smaller than long-run responses. In contrast, when $\bar{\pi}_{12} < 0$, an agent experiences more severe historical restraints following a change in climate than following a transient weather shock. When actions depend on scarce resources, actions can be more extreme when they are maintained for only a short period of time. The $\omega > 1$ captures how responses to transient shocks understate historical restraints in this case. Consistent with conjectures in Fisher et al. (2012) and Blanc and Schlenker (2017), short-run responses are larger than long-run responses.

The wedges introduced by Ω and ω conflict, making it impossible to sign the bias in general. However, we can make progress in two special cases. First, when $\bar{\pi}_{12} = 0$, both wedges vanish. In this case, the fixed effects estimator exactly recovers the effect of climate. Second, when $\beta = 0$, 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).²⁰

Now consider the following distributed lag regression, which matches most literature in not controlling for lagged actions:²¹

$$A_{jt} = \alpha_j + \sum_{i=0}^{I+1} \Gamma_{w_{t-i}} w_{j(t-i)} + \sum_{i=0}^{I+1} \Gamma_{f_{t-i}} f_{j(t-i)} + \eta_{jt}, \quad (11)$$

where $I \geq 0$. The following proposition relates the estimated coefficients to the effect of climate change.²²

Proposition 2. *Let either Assumption 1 or 2 hold, and let $(A_{j(t-1)} - \bar{A})^2$ be small for all observations. Then*

$$\lim_{I \rightarrow \infty} \sum_{i=0}^I \left[\hat{\Gamma}_{w_{t-i}} + \hat{\Gamma}_{f_{t-i}} \right] = \tilde{\omega} \left(\frac{d\bar{A}}{dC} + \beta \bar{\pi}_{12} \Omega \right),$$

²⁰The wedge introduced by preparatory actions also vanishes if there is no ex-post adaptation, but this is an artifact of modeling forecasts as existing only one period ahead. In this environment, there are no time t shocks that affect expectations of time $t + 2$ weather. If there were longer-horizon forecasts, then time t shocks could affect those expectations and thereby induce preparation for ex-ante adaptation anticipated to be undertaken at time $t + 1$.

²¹This omission is driven in part by concern for Nickell (1981) omitted variables bias.

²²The proof also shows that, beyond the first lag, the coefficients alternate signs as the lags increase if and only if $\bar{\pi}_{12} < 0$. The appendix shows that $\bar{p}_{12} = 0$ is not sufficient to recover the effects of climate if the regression does not control for forecasts and there is scope for ex-ante adaptation.

where $\beta\bar{\pi}_{12} > 0$ implies $\tilde{\omega} \in (\omega, 1)$, $\beta\bar{\pi}_{12} < 0$ implies $\tilde{\omega} \in (1, \omega)$, and $\beta\bar{\pi}_{12} = 0$ implies $\tilde{\omega} = 1$, with ω and Ω from Proposition 1. If $\bar{\pi}_{12} = 0$, then $\hat{\Gamma}_{w_{t-i}} = 0$ for $i > 1$ and $\hat{\Gamma}_{f_{t-i}} = 0$ for $i > 0$. If $\beta = 0$, then $\hat{\Gamma}_{f_{t-i}} = 0$ for all $i \geq 0$.

Proof. See appendix. □

This estimator is subject to the same bias from preparatory actions, but by using a long history of transient shocks, it reduces the bias introduced by historical restraints. If the latter bias is the dominant one, then this estimator may reduce the overall bias in estimated effects on actions. Further, this estimator recovers the effects of climate in a new case: when agents are myopic. Myopic agents are never subject to the bias induced by preparatory actions, and we now lose the bias induced by historical restraints because 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.

4.2 Estimating Effects on Payoffs

We now consider the possibility of recovering effects on payoffs from observations of payoffs and weather. For instance, empirical research studies how variation in weather affects agricultural profits (e.g., Deschênes and Greenstone, 2007) or affects macroeconomic variables such as gross output or income that are potentially related to payoffs (e.g., Dell et al., 2012; Burke et al., 2015; Deryugina and Hsiang, 2017; Colacito et al., 2019).

The class of payoff functions defined by the following assumption will yield especially interesting results:

Assumption 3. $\pi_2(A_t, A_{t-1}, w_t, w_{t-1}) = K\pi_1(A_t, A_{t-1}, w_t, w_{t-1})$ if $A_{t-1} = A_t$, for $K \neq -\beta$.

Consider a few members of this class. First, adjustment cost models yield $K = 0$: if $\pi = g(A_t, (A_t - A_{t-1})^z, w_t, w_{t-1})$ for $z > 1$, then $\pi_2 = z(A_t - A_{t-1})^{z-1}g_2(A_t, (A_t - A_{t-1})^z, w_t, w_{t-1})$ and thus is equal to 0 when $A_t = A_{t-1}$. Second, a model in which the returns to resource extraction decline in previous extraction can yield $K = -1$: if $\pi = g(A_t/A_{t-1}, w_t, w_{t-1})$, then $\pi_1 = g_1(A_t/A_{t-1}, w_t, w_{t-1})/A_{t-1}$ and $\pi_2 = -A_t g_1(A_t/A_{t-1}, w_t, w_{t-1})/A_{t-1}^2$. Third, a model in which ex-post adaptation and ex-ante adaptation form a constant elasticity of substitution (CES) aggregate with distribution parameter κ yields $K = (1 - \kappa)/\kappa$: $\pi = g(h(A_t, A_{t-1}), w_t, w_{t-1})$ where $h(A_t, A_{t-1}) = (\kappa A_t^\sigma + (1 - \kappa)A_{t-1}^\sigma)^{1/\sigma}$ for $\sigma < 1, \neq 0$ and $h(A_t, A_{t-1}) \rightarrow A_t^\kappa A_{t-1}^{1-\kappa}$ as $\sigma \rightarrow 0$. Finally, a model without dynamic linkages has $\pi_2(\cdot, \cdot, \cdot, \cdot) = 0$ and thus $K = 0$.

Empirical researchers hope to recover (6) from time series variation in weather. They will not generally observe the full set of actions available to agents or firms. As a result, empirical researchers may estimate the following regression:

$$\pi_{jt} = \alpha_j + \sum_{i=0}^{I+1} \theta_{w_{t-i}} w_{j(t-i)} + \sum_{i=0}^{I+1} \theta_{f_{t-i}} f_{j(t-i)} + \eta_{jt}, \quad (12)$$

where I again label units as j , α_j is a fixed effect for agent j , and η_{jt} is an error term (see footnote 19). We are interested in the vector of coefficients θ . As before, I use a hat to denote the probability limit of each coefficient.

Proposition 3. *Let Assumption 1 hold, or let Assumption 2 hold with the ϵ and ν normally distributed. Also let $(A_{j(t-1)} - \bar{A})^2$ and $(A_{jt} - \bar{A})^2$ be small for all observations and let each agent's average actions be \bar{A} .*

1. *If $\bar{\pi}_{12} = 0$ and $I > 1$, then $\lim_{s \rightarrow \infty} dE_0[\pi_s]/dC = \hat{\theta}_{w_t} + \hat{\theta}_{w_{t-1}} + \hat{\theta}_{w_{t-2}} + \hat{\theta}_{f_t} + \hat{\theta}_{f_{t-1}}$ and all other coefficients are equal to 0.*
2. *If $\beta = 0$, then $\lim_{s \rightarrow \infty} dE_0[\pi_s]/dC = \lim_{I \rightarrow \infty} \left[\sum_{i=0}^I \hat{\theta}_{w_{t-i}} + \sum_{i=0}^I \hat{\theta}_{f_{t-i}} \right]$.*
3. $\lim_{s \rightarrow \infty} dE_0[\pi_s]/dC = \lim_{\beta \rightarrow 1} \lim_{I \rightarrow \infty} \left[\sum_{i=0}^I \hat{\theta}_{w_{t-i}} + \sum_{i=0}^I \hat{\theta}_{f_{t-i}} \right]$.
4. *If Assumption 3 holds and $I \geq 0$, then $\lim_{s \rightarrow \infty} dE_0[\pi_s]/dC = \hat{\theta}_{w_t} + \hat{\theta}_{w_{t-1}}$ and all other coefficients are equal to zero.*

Proof. See appendix. □

The proposition describes four cases in which we can recover the effect of climate from time series variation in weather (and the appendix further describes cases in which we can unambiguously bound the effect). The first two cases follow directly from the analysis in Section 4.1. There we saw that we can recover the effect of climate on actions if either $\bar{\pi}_{12} = 0$ or $\beta = 0$. In the former case, we can recover the effect on current actions from the coefficient on weather, its lag, and forecasts, and the first result in Proposition 3 follows from recognizing that we need to recover effects on both current and lagged actions and that the coefficients on weather and its lag also capture the direct effects of weather in equation (6). If $\beta = 0$, we recover effects on actions only as the lags become very long, in which case we also recover effects on payoffs.

The other two cases are ones in which we do not need to recover the effect of climate on actions. The proof shows that the bias from estimating the effect of climate on payoffs from the combination of infinite lags is proportional to $\beta \bar{\pi}_{12} (\bar{\pi}_1 + \bar{\pi}_2)$. The bias vanishes as $\beta \rightarrow 1$ because agents' responses equalize the marginal value of past and current actions, without discounting the former. Alternately, Assumption 3 and equation (2) imply that $\bar{\pi}_2 = \bar{\pi}_1 = 0$: an optimizing agent sets the marginal benefit of actions to zero around a steady state. In this case, the consequences of marginal climate change are independent of changes in actions and the estimated coefficients do not include any effects of weather or forecasts on actions. Summing $\hat{\theta}_{w_t}$ and $\hat{\theta}_{w_{t-1}}$ now captures only the direct effects of weather and fully captures the effects of climate on payoffs. We can test whether Assumption 3 holds by examining the magnitude of the coefficient on forecasts: because forecasts matter for current payoffs




Figure 3: Per-period payoffs π_t against previous actions A_{t-1} . Each solid curve holds time t actions and weather fixed. The dotted curve gives the effect of time $t - 1$ actions on myopically maximized payoffs.

only through their effects on actions, they cannot affect these payoffs if Assumption 3 indeed holds and agents are near a steady state.^{23,24}

Assumption 3 leads us to the same point as envelope theorem intuition proposed in previous literature. But Section 3 explained why these arguments fail. So what does Assumption 3 do? Figure 3 plots payoffs as a function of previous actions, conditional on current actions and weather. The dotted curve gives the effect of previous actions on myopically optimized payoffs. Point (1) identifies the previous action that maximizes current payoffs. Around this point, previous actions do not have first-order effects on current payoffs. Why is that important? The dashed curves in Figure 1 reflected how past actions can shift current payoffs. If these past actions do not have first-order effects, then the dashed curves are not shifted out and point (c) converges to point (b) in Figure 1. Further, if past actions do not have first-order effects on current payoffs, then current actions do not have first-order effects on future payoffs. In that case, the myopically optimal action is also dynamically optimal and point (d) converges to point (b) in Figure 1. Assumption 3 ensures that point (1) occurs at the steady-state action. Therefore, if Assumption 3 holds and agents are near a steady state, then points (c) and (d) converge to point (b) in Figure 1 and the treatment effect of a transient weather shock indeed recovers the effect of permanently changing the weather.²⁵

²³Much literature has studied dependent variables such as crop yields (e.g., Schlenker and Roberts, 2009), mortality (e.g., Deschênes and Moretti, 2009; Deschênes and Greenstone, 2011), and health (e.g., Deschenes, 2014) that are functions of actions but are not payoff functions. If we consider recovering the effects of climate on such dependent variables from a fixed effects regression on weather, then the final two parts of Proposition 3 no longer apply because the Euler equation (1) holds only for payoffs, not for other functions of actions.

²⁴Empirical researchers have usually not controlled for forecasts. The appendix shows that such regressions can still recover the full effect of climate on payoffs if either Assumption 3 holds or agents are myopic. However, the other cases now also require the absence of ex-ante adaptation.

²⁵As described earlier, one of the special cases of Assumption 3 is a model with no dynamic linkages

Proposition 3 assumed that each agent's average actions are \bar{A} .²⁶ The following corollary establishes how relaxing this assumption changes the results.

Corollary 4. *Let the conditions given in Proposition 3 hold, except let each agent's average actions be different from \bar{A} . In addition, let at least one of $\bar{\pi}_{13}$, $\bar{\pi}_{23}$, $\bar{\pi}_{14}$, or $\bar{\pi}_{24}$ be strictly positive. Then, in each part of Proposition 3, $\lim_{s \rightarrow \infty} dE_0[\pi_s]/dC$ is strictly less (greater) than the indicated combination of coefficients if and only if \bar{A} is strictly less (greater) than each agent's average actions.*

Proof. See appendix. □

The corollary establishes that the special cases that formerly sufficed to identify climate impacts from weather impacts now merely bound the effect of climate on payoffs. In particular, we obtain an upper bound if agents are approaching their steady-state actions from above and a lower bound otherwise. Intuitively, if climate shifts the steady-state action farther from the agent's current action, then weather shocks incorporate transition costs that vanish from the effect of climate on long-run payoffs.

5 Structurally Estimating Climate Impacts in U.S. Agriculture

I have thus far explored the potential for reduced-form estimates to recover climate impacts. I now show how to recover the combinations of structural parameters necessary to calculate climate impacts. I use these parameters to disentangle weather effects from adaptation and to sign the bias in the adaptation estimate. 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 specify or recover every underlying structural parameter in order to undertake the calculations of interest.

I demonstrate this new approach by extending the seminal analysis of agricultural impacts from Deschênes and Greenstone (2007). In order to be consistent with common regression specifications, generalize the foregoing analysis to allow for K types of weather variables, which can be correlated with each other. Let there be M actions chosen in each period, so that time t payoffs are now $\pi(\mathbf{A}_t, \mathbf{A}_{t-1}, \mathbf{w}_t, \mathbf{w}_{t-1}) : \mathbb{R}^M \times \mathbb{R}^M \times \mathbb{R}^K \times \mathbb{R}^K \rightarrow \mathbb{R}$, with bold script indicating vectors. Superscripts will indicate elements of these vectors. The following assumption is useful for the structural calculations:

$(\pi_2(\cdot, \cdot, \cdot, \cdot) = 0)$, in which case the agent solves a series of independent, static decision problems. Appeals to the envelope theorem therefore can end up with the correct result in the types of static settings discussed by previous literature (Hsiang, 2016; Deryugina and Hsiang, 2017).

²⁶Sorkin (2015) imposes an analogous restriction when relating short-run and long-run variation.

Assumption 4. *Either (i) there exist $g : \mathbb{R}^M \rightarrow \mathbb{R}$, $h : \mathbb{R}^M \rightarrow \mathbb{R}$, and $\pi^0 : \mathbb{R} \times \mathbb{R} \times \mathbb{R}^K \rightarrow \mathbb{R}$ such that $\pi(\mathbf{A}_t, \mathbf{A}_{t-1}, \mathbf{w}_t, \mathbf{w}_{t-1}) = \pi^0(g(\mathbf{A}_t), h(\mathbf{A}_{t-1}), \mathbf{w}_t)$, or (ii) there exist $\pi^k : \mathbb{R} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ for $k \in \{1, \dots, K\}$ and $\pi^{K+1} : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}$ such that $M \geq K$ and*

$$\pi(\mathbf{A}_t, \mathbf{A}_{t-1}, \mathbf{w}_t, \mathbf{w}_{t-1}) = \sum_{k=1}^K \pi^k(A_t^k, A_{t-1}^k, w_t^k) + \pi^{K+1}(\mathbf{A}_t^{\sim k}, \mathbf{A}_{t-1}^{\sim k}),$$

where $\mathbf{A}_t^{\sim k}$ indicates the $(M - K)$ -dimensional vector of actions A_t^{K+1} through A_t^M .

This assumption does two things. First, it rules out delayed effects of weather, which is plausible in the below application to annual agricultural data. Some assumption about delayed effects is necessary for identification of structural parameters. Second, it requires either that the vector of actions be reducible to a composite action (trivially true for $M = 1$) or that payoffs be separable in the dimensions of weather. Either restriction ensures that the terms controlling whether actions are intertemporal substitutes or complements are scalar.

Consider the following regression:

$$\pi_{ct} = \alpha_c + \psi_{rt} + \sum_{k=1}^K [\Phi_{w_{t-2}}^k w_{c(t-2)}^k + \Phi_{w_{t-1}}^k w_{c(t-1)}^k + \Phi_{w_t}^k w_{ct}^k + \Phi_{w_{t+1}}^k w_{c(t+1)}^k] + \delta_{ct}, \quad (13)$$

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,²⁷ and superscript k indexes weather variables of interest. The following lemma expresses the coefficients in terms of model primitives.²⁸

Lemma 3. *Let Assumption 4 and the conditions of Proposition 3 hold, and let ϵ_t be uncor-*

²⁷In the preferred specification, the regions are USDA Farm Resource Regions (see also Deschênes and Greenstone, 2012). The appendix provides further details, reports the variance explained by the weather variables (see Fisher et al., 2012), and 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).

²⁸The lemma requires that weather be serially uncorrelated. This assumption seems an acceptable starting point: over all U.S. counties from 1972 to 2017, the correlation between locally demeaned growing season degree days and its lag is 0.10, the correlation between locally demeaned extreme growing season degree days and its lag is 0.074, and the correlation between locally demeaned growing season precipitation and its lag is -0.029.

related with ν_t . Then:

$$\begin{aligned}\hat{\Phi}_{w_{t+1}}^k &= -\beta \bar{\pi}_2^k \hat{\Gamma}_3^k \frac{(\tau^k)^2}{(\tau^k)^2 + (\sigma^k)^2}, \\ \hat{\Phi}_{w_t}^k &= \bar{\pi}_3^k - \beta \bar{\pi}_2^k \hat{\Gamma}_1^k + \left(1 - \beta \frac{\bar{\pi}_{12}^k}{\chi_2^k}\right) \bar{\pi}_2^k \frac{(\tau^k)^2}{(\tau^k)^2 + (\sigma^k)^2} \hat{\Gamma}_3^k, \\ \hat{\Phi}_{w_{t-1}}^k &= \left(1 - \beta \frac{\bar{\pi}_{12}^k}{\chi_2^k}\right) \bar{\pi}_2^k \hat{\Gamma}_1^k + \left(1 - \beta \frac{\bar{\pi}_{12}^k}{\chi_2^k}\right) \frac{\bar{\pi}_{12}^k}{\chi_2^k} \bar{\pi}_2^k \frac{(\tau^k)^2}{(\tau^k)^2 + (\sigma^k)^2} \hat{\Gamma}_3^k, \\ \hat{\Phi}_{w_{t-2}}^k &= \frac{\bar{\pi}_{12}^k}{\chi_2^k} \left(1 - \beta \frac{\bar{\pi}_{12}^k}{\chi_2^k}\right) \bar{\pi}_2^k \hat{\Gamma}_1^k + \left(\frac{\bar{\pi}_{12}^k}{\chi_2^k}\right)^2 \left(1 - \beta \frac{\bar{\pi}_{12}^k}{\chi_2^k}\right) \bar{\pi}_2^k \frac{(\tau^k)^2}{(\tau^k)^2 + (\sigma^k)^2} \hat{\Gamma}_3^k.\end{aligned}$$

If case (i) of Assumption 4 holds, then we replace the k superscripts on the right-hand side with 0.

Proof. See appendix. □

The $\hat{\Gamma}$ were defined in regression (11) and analyzed in Proposition 1. The proof of Proposition 1 defines these in terms of model primitives. They here gain a superscript k to indicate the corresponding dimension of weather. Solving the system of equations, we find:

$$\begin{aligned}\frac{\bar{\pi}_{12}^k}{\chi_2^k} &= \frac{\hat{\Phi}_{w_{t-2}}^k}{\hat{\Phi}_{w_{t-1}}^k}, \\ \bar{\pi}_2^k \hat{\Gamma}_3^k \frac{(\tau^k)^2}{(\tau^k)^2 + (\sigma^k)^2} &= -\frac{\hat{\Phi}_{w_{t+1}}^k}{\beta}, \\ \bar{\pi}_2^k \hat{\Gamma}_1^k &= \frac{\hat{\Phi}_{w_{t-1}}^k - \left(1 - \beta \frac{\bar{\pi}_{12}^k}{\chi_2^k}\right) \frac{\bar{\pi}_{12}^k}{\chi_2^k} \bar{\pi}_2^k \frac{(\tau^k)^2}{(\tau^k)^2 + (\sigma^k)^2} \hat{\Gamma}_3^k}{1 - \beta \frac{\bar{\pi}_{12}^k}{\chi_2^k}}, \\ \bar{\pi}_3^k &= \hat{\Phi}_{w_t}^k + \beta \bar{\pi}_2^k \hat{\Gamma}_1^k - \left(1 - \beta \frac{\bar{\pi}_{12}^k}{\chi_2^k}\right) \bar{\pi}_2^k \frac{(\tau^k)^2}{(\tau^k)^2 + (\sigma^k)^2} \hat{\Gamma}_3^k.\end{aligned}$$

$\bar{\pi}_{12}^k/\chi_2^k$ is identified from the first and second lags of weather. Recalling that χ_2^k must be positive for a saddle-path stable steady-state, the sign of this term matches the sign of $\bar{\pi}_{12}^k$. The lead of weather identifies $\bar{\pi}_2^k \hat{\Gamma}_3^k \frac{(\tau^k)^2}{(\tau^k)^2 + (\sigma^k)^2}$, which the proof of Proposition 1 connects to ex-ante adaptation.²⁹ Given these two terms, the residual effects of lagged weather identify $\bar{\pi}_2^k \hat{\Gamma}_1^k$, which the proof of Proposition 1 connects to ex-post adaptation. Finally, the residual effects of contemporary weather identify the direct effects $\bar{\pi}_3^k$. Intuitively, lagged weather affects current payoffs only through actions, so we can identify ex-post adaptation from

²⁹I calibrate β to the annual discount rate of 34% obtained in Duquette et al. (2012). The appendix shows that results are not sensitive to lower discount rates.

lagged weather and can identify intertemporal links between actions by comparing lags of weather. And once we have identified the scope of adaptation, we can recover the direct effects of weather from the response of payoffs to contemporary weather.

Equation (7) shows that calculating climate impacts requires $\bar{\pi}_3$ and $(1 - \beta)\bar{\pi}_2 d\bar{A}/dC$. The above steps recover $\bar{\pi}_3$ directly. From equation (10),

$$\frac{d\bar{A}}{dC} = \frac{1}{\omega} \left[\hat{\Gamma}_1 + \hat{\Gamma}_3 \right] - \beta\bar{\pi}_{12}\Omega,$$

with $\omega > 1$ if and only if $\bar{\pi}_{12} < 0$. If $\bar{\pi}_{12} \approx 0$, then

$$\frac{d\bar{A}}{dC} \approx \hat{\Gamma}_1 + \hat{\Gamma}_3.$$

I calculate the effects of ex-post adaptation using $(1 - \beta)\bar{\pi}_2\hat{\Gamma}_1$ and the effects of ex-ante adaptation using $(1 - \beta)\bar{\pi}_2\hat{\Gamma}_3$.³⁰ As described below, I use the estimated sign of $\bar{\pi}_{12}^k$ to convert these estimates into bounds.

For comparison, I also undertake two reduced-form calculations of the effects of climate change. Either can be justified from an analogue of Proposition 3. A first calculation estimates (13) without any leads or lags on the right-hand side and multiplies each weather variable's coefficient by the projected change in that weather variable. This calculation recovers the theory-implied effects of climate if Assumptions 3 and 4 hold. It matches the calculations undertaken in previous literature. A second calculation estimates (13) without the second lag on the right-hand side and multiplies the sum of each weather index's three coefficients by the projected change in that weather variable. This calculation recovers the theory-implied effects of climate if $\bar{\pi}_{12} = 0$ and Assumption 4 holds.

The appendix describes the data, details sample construction, and undertakes robustness checks. 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 2012. 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), 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 meridian, which are less likely to be irrigated (Schlenker et al., 2005; Fisher et al., 2012).

³⁰These calculations set $(\tau^k)^2/((\tau^k)^2 + (\sigma^k)^2)$ equal to 1. This fraction reflects the fraction of the variation in weather that is already realized one period ahead (i.e., that is reflected in forecasts). Estimated ex-ante adaptation is biased towards zero if the fraction is in fact less than 1. Replacing the lead of realized weather in regression (13) with forecasts would eliminate this bias.

Table 1: Top panel: Estimated coefficients and standard errors from regression (13). Bottom panel: Theory-implied structural parameters, reported as the median and lower/upper quartiles of the distribution implied by the reduced-form estimates.

	GDD	Extreme GDD	Precip
<i>Reduced-Form Coefficients</i>			
Current	18 (9.1)	-98 (54)	-7.9 (3.9)
Lag 1	5.4 (6.2)	-60 (22)	-7.1 (3.9)
Lag 2	-15 (9.9)	32 (23)	0.51 (3.8)
Lead	-3.8 (4.1)	-1.4 (20)	1.3 (1.5)
<i>Theory-Implied Parameters</i>			
$\bar{\pi}_3$	15 (8.9,22)	-130 (-160,-110)	-12 (-16,-7.8)
$\bar{\pi}_2 \hat{\Gamma}_1$	13 (3.4,28)	-46 (-60,-31)	-7.6 (-11,-5)
$\bar{\pi}_2 \hat{\Gamma}_3 \frac{\tau^2}{\tau^2 + \sigma^2}$	5.0 (1.3,8.8)	1.9 (-16,20)	-1.8 (-3.2,-0.43)
$\bar{\pi}_{12}/\chi_2$	-1.7 (-3.4,-0.41)	-0.53 (-0.74,-0.31)	-0.065 (-0.46,0.31)

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 13944 county-year observations and 37 state observations. Profits in thous. year 2002 dollars, GDD in °C-days, and precip in mm.

The top panel of Table 1 reports the reduced-form coefficients from regression (13). Profits increase in same-year and previous-year growing degree days between 10°C and 29°C (“GDD”), but profits decrease in same-year and previous-year growing degree days above 29°C (“Extreme GDD”). The signs of the central estimates alternate from the first to the second lag. The lead of weather does not have a statistically significant effect on profits. Additional same-year and previous-year precipitation reduces profits.

The lower panel of Table 1 reports the medians and, in parentheses, lower and upper quartiles for the theory-implied structural parameters.³¹ The signs of the direct terms $\bar{\pi}_3$ and the ex-post adaptation terms $\bar{\pi}_2\hat{\Gamma}_1$ are consistent with the signs of same-year and previous-year impacts on profits. The nonzero ex-post adaptation term suggests that Assumption 3 does not hold in this application, which motivates a structural approach to recovering climate impacts. The ex-ante adaptation terms are small and noisy, as were the coefficients on the leads of weather.

The final row of Table 1 reports $\bar{\pi}_{12}/\chi_2$. We have a case of intertemporal substitutes (complements) if this term is negative (positive). The median values are all negative, and even the 75th percentile is negative for the growing degree day variables.³² The negative $\bar{\pi}_{12}$ are identified off the result that lagged weather affects payoffs much like current weather but the second lag of weather has opposite effects to current weather. Within the economic model, these opposite effects imply that adaptive actions taken two years ago increase current payoffs by constraining the actions taken last year. Finding $\bar{\pi}_{12} < 0$ 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.³³

Table 2 reports the projected effects of climate change. It uses the RCP 4.5 trajectory of stabilized emissions from 21 downscaled CMIP5 models. In this scenario, global mean surface temperature increases by around 2 degrees Celsius over the century. The top panel reports the two reduced-form calculations. Both approaches project substantial costs from climate change. The projected increase in conventional growing degree days is estimated to

³¹The lower panel does not report means and standard errors because the distributions can be skewed due to $\bar{\pi}_{12}/\chi_2$ being the ratio of two reduced-form coefficients. The appendix reports standard errors from method of moments estimators of the structural parameters and reports results for the 10th and 90th percentiles.

³²Reassuringly, these estimates are consistent with $|\bar{\pi}_{12}| < \chi_2$, which the proof of Lemma 2 shows follows from saddle-path stability. One might consider estimating a geometric lag structure in regression (13), with the geometric term equal to $\bar{\pi}_{12}/\chi_2$. I find that coefficient estimates can be unstable when adding longer lags, likely due to collinearity problems that commonly arise with lagged predictors. When using three lags, the estimated $\bar{\pi}_{12}/\chi_2$ for extreme growing degree days is extremely noisy, the estimated $\bar{\pi}_{12}/\chi_2$ for conventional growing degree days in a one- (two-) step GMM estimator is -1.18 (-0.56) with a standard error of 0.43 (0.34), and the estimated $\bar{\pi}_{12}/\chi_2$ for precipitation in a one- (two-) step GMM estimator is -1.47 (-0.17) with a standard error of 0.56 (0.31).

³³Huang and Moore (2019) show that U.S. farmers do adjust crop acreage in response to pre-planting precipitation.

increase agricultural profits, but the projected increase in extreme growing degree days is projected to reduce profits to a greater degree. Projected costs are 50% greater under the assumption that $\bar{\pi}_{12} = 0$ than under Assumption 3. However, these numbers are meaningless if Assumption 3 does not hold or $\bar{\pi}_{12} \neq 0$, and we have already seen that Assumption 3 does not appear to hold and that $\bar{\pi}_{12} < 0$.

The lower panel reports the new, theory-based estimates of climate impacts. It decomposes the effects of climate change into direct effects (driven by $\bar{\pi}_3$), ex-post adaptation (driven by $\bar{\pi}_{13}$ via $\bar{\pi}_2 \hat{\Gamma}_1$), and ex-ante adaptation (driven by $\bar{\pi}_{23}$ via $\bar{\pi}_2 \hat{\Gamma}_3$). The median direct effects of projected changes in each measure of growing degree days are smaller than either of the reduced-form estimates, but the median combined direct effect is in between the two reduced-form estimates.

Ex-post adaptation increases the benefits from conventional growing degree days but increases the costs from extreme growing degree days. From equation (6), adaptive changes in actions affect steady-state payoffs as $\bar{\pi}_1 + \bar{\pi}_2$, and from equation (2), $\bar{\pi}_1 + \bar{\pi}_2$ is opposite in sign to $\bar{\pi}_1$. As described following equation (7), actions that provide short-run benefits (costs) in exchange for long-run costs (benefits) have negative (positive) effects on steady-state payoffs when agents are not perfectly patient. Adaptation to increases in conventional growing degree days therefore imposes short-run costs, but agents reap the benefits from past adaptation in the long run. In contrast, adaptation to increases in extreme growing degree days provides short-run benefits, but agents reap these benefits only while paying the larger costs of past adaptation, which dominate in the long run. Ex-ante adaptation is not clearly important.³⁴ Accounting for adaptation, projected changes in conventional growing degree days increase profits by 54% in the median estimate and projected changes in extreme growing degree days eliminate profits in the median estimate. The median total effect of climate change is a 50% reduction in profits.

The median effect of climate happens to be somewhat close to the value that would be obtained by applying previous literature's methods (which are identical to applying Assumption 3). However, the interpretation is quite different. The structural estimates approximate $d\bar{A}/dC$ by setting $\bar{\pi}_{12} = 0$, which implies that $\omega = 1$ and that Ω is irrelevant in equation (10). In contrast to reduced-form approaches, these structural calculations have a clear interpretation even if that assumption is violated. We know from Proposition 1 that the bias from $\bar{\pi}_{12} \neq 0$ reflects preparatory actions (through Ω) and historical restraints (through ω). In the present context, it is reasonable to assume that the wedge induced by preparatory actions is small relative to the wedge induced by historical restraints, an intuition reinforced by the small effects of ex-ante adaptation. Therefore $\bar{\pi}_{12} > 0$ implies that the present calculations

³⁴Recall, however, that the estimated effects of ex-ante adaptation are biased towards zero (see footnote 30). The appendix reports that ex-ante adaptation to extreme growing degree days becomes clearer if we omit the earlier years from the sample, which is consistent with the reduced skill and availability of seasonal forecasts prior to the mid-1990s (see Klemm and McPherson, 2017). Takle et al. (2013) and Klemm and McPherson (2017) describe the various seasonal forecasts of interest to agriculture.

Table 2: 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 estimates and standard errors. The theory-implied estimates report the median and lower/upper quartiles.

	GDD	Extreme GDD	Precip	Combined
<i>Reduced-Form</i>				
Using Assumption 3	51 (25)	-97 (37)	-1.0 (0.4)	-47 (24)
Using $\bar{\pi}_{12} = 0$	44 (29)	-1.1e+02 (35)	-1.6 (0.66)	-71 (28)
<i>Theory-Implied</i>				
Direct Effects	38 (22,54)	-93 (-1.1e+02,-74)	-1.2 (-1.7,-0.83)	-56 (-71,-41)
Ex-Post Adaptation	8.3 (2.1,17)	-8.0 (-11,-5.4)	-0.21 (-0.29,-0.13)	0.56 (-5.7,9.2)
Ex-Ante Adaptation	3.2 (0.84,5.5)	0.33 (-2.8,3.4)	-0.048 (-0.085,-0.012)	3.4 (0.026,6.8)
Total	54 (32,74)	-1e+02 (-1.2e+02,-79)	-1.5 (-2,-1)	-50 (-69,-31)

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 13944 county-year observations and 37 state observations.

underestimate adaptation to climate (because $\omega < 1$) and $\bar{\pi}_{12} < 0$ implies that the present calculations overestimate adaptation to climate (because $\omega > 1$).

We saw in Table 1 that $\bar{\pi}_{12} < 0$, implying that we observe more adaptation to short-run weather shocks than would occur in response to long-run changes in climate. We can therefore bound the effects of climate by the estimated total effects that include projected adaptation and by the estimated direct effects that exclude adaptation. If I were following the conventional approach in the weather-climate literature, I would undertake calculations like those relying on Assumption 3 and apply intuition based on Le Châtelier's principle to conclude that climate change reduces profits by 0–47% in the median estimates. Instead, here the median estimates and the finding that $\bar{\pi}_{12} < 0$ imply that climate change reduces agricultural profits by 50–56% (\$20–22 billion annually at the year 2017 price level), with the 75th percentile estimates yielding a reduction in profits of 31–41% and the 90th percentile estimates (see appendix) yielding a reduction in profits of 9–28%.

6 Potential Extensions

I have demonstrated how to estimate the effects of climate change from time series variation in weather. I conclude by discussing the primary restrictions in the present setting and describing other aspects of climate change that should be the subject of future analysis.

The theoretical model is fairly general. The three notable restrictions are that past weather and actions can directly affect payoffs with only a one-period lag (although they do indirectly affect payoffs arbitrarily far into the future) and that agents have access to specialized forecasts only one period in advance of a realized weather outcome. Longer-run delayed weather impacts are probably important to some applications, but they would not substantially change the theoretical results. In particular, distributed lag models will recover the effects of climate in exactly the same cases as analyzed here. Allowing for longer-horizon forecasts also does not change the theoretical results unless actions can have direct effects over those horizons.

Allowing actions to have longer-run direct consequences can have interesting consequences. Such an extension is attractive if we interpret actions as the choice of capital stock and we want to study capital that can be built only with a lag of more than one period. If longer-run forecasts are available, then the present results extend in a natural way, implying that it is important to control for these longer-run forecasts. But if weather forecasts do not exist over the whole horizon over which today's actions will directly affect payoffs, then empirical researchers may be unable to estimate the full effect of climate: changing the climate can lead agents to undertake actions that pay off only in the distant future, but observable variation in forecasts will not identify this adaptation margin. In the empirical application of Section 5, the existence of such actions would imply that long-run adaptation could be greater than short-run adaptation even though I estimate $\bar{\pi}_{12} < 0$ from the actions that vary in the data.

The present setting successfully captures the distinction between transient and permanent changes in weather. Future work should consider other aspects of climate change. First, global climate change differs from weather shocks not only in its temporal structure but also in its spatial structure. A change in global climate affects weather in every location and thus will have general equilibrium consequences. In the empirical application, general equilibrium channels change the prices of land and crops. The present setting has followed most empirical work in abstracting from such effects, but some recent empirical work has begun exploring the implications of changing the weather in many locations simultaneously (e.g., Costinot et al., 2016; Gouel and Laborde, 2018; Dingel et al., 2019). Future work should extend the present setting to account for general equilibrium effects.

Second, the present analysis has held the payoff function fixed over time. However, climate change should induce innovations that alter how weather affects payoffs, and many such innovations will arise even in the absence of climate change. Some types of innovation can be interpreted as actions within the present framework, but the potential for future

innovation may be inherently unobservable. Historical studies have begun exploring the interaction between climate and agricultural innovation (e.g., Olmstead and Rhode, 2008, 2011; Roberts and Schlenker, 2011; Bleakley and Hong, 2017). Future work should consider approaches to bounding the scope for future innovation.

Third, the present analysis has considered only marginal changes in climate, but climate change over the next century is likely to be nonmarginal.³⁵ The present analysis implies that the marginal effect of climate will not generally be constant. Some recent work (Mérel and Gammans, 2019) explores the types of variation captured by quadratic regression terms. Future work should explore whether nonlinear specifications might inform estimates of the impacts from nonmarginal climate change in a dynamic setting such as the present one.

Fourth, the present analysis assumes that actions can be adjusted continuously. In the presence of fixed costs, an agent may choose to change an action only when the agent expects a change in weather to endure, and in the presence of constraints imposed by policy, actions may not respond smoothly to weather shocks, changing the interpretation of reduced-form coefficients. Future work should explore the conditions under which aggregating over many agents' fixed-cost decisions makes actions appear continuous. Future work should also explore whether responses to weather events of varying durations can identify how fixed-cost actions respond to a change in climate and whether responses to weather events of varying magnitudes can identify the role of policy constraints.

Finally, the present analysis has focused on identifying the long-run consequences of climate change, abstracting from the transition costs induced by climate change. In this regard, the present analysis matches the calculations undertaken by nearly all empirical work but omits a potentially critical aspect of climate change (see Quiggin and Horowitz, 1999, 2003; Kelly et al., 2005). Future work should consider whether imposing stronger assumptions on the decision-making environment can allow for credible simulation of counterfactual climate trajectories and thereby estimate transition costs.

References

Anderson, Michael L. (2015) "As the wind blows: The effects of long-term exposure to air pollution on mortality," Working Paper 21578, National Bureau of Economic Research.

Auffhammer, Maximilian (2018a) "Climate adaptive response estimation: Short and long run

³⁵Estimating the consequences of nonmarginal climate change is critical to the damage functions required by climate-economy integrated assessment models (see Nordhaus, 2013). However, there is an argument that the consequences of marginal climate change might be especially policy relevant: if we accept climate scientists' views that the potentially nonquantifiable risks imposed by nonmarginal climate change are likely to exceed the cost of avoiding them, then the effects of marginal climate change become critical to policy choices.

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.

Auffhammer, Maximilian and Anin Aroonruengsawat (2011) “Simulating the impacts of climate change, prices and population on California’s residential electricity consumption,” *Climatic Change*, Vol. 109, No. 1, pp. 191–210.

Bakkensen, Laura and Lint Barrage (2018) “Climate shocks, cyclones, and economic growth: Bridging the micro-macro gap,” Working Paper 24893, National Bureau of Economic Research.

Barreca, Alan I., Matthew Neidell, and Nicholas J. Sanders (2017) “Long-run pollution exposure and adult mortality: Evidence from the Acid Rain Program,” Working Paper 23524, National Bureau of Economic Research.

Bishop, Kelly C., Jonathan D. Ketcham, and Nicolai V. Kuminoff (2018) “Hazed and confused: The effect of air pollution on dementia,” Working Paper 24970, National Bureau of Economic Research.

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.

Bleakley, Hoyt and Sok Chul Hong (2017) “Adapting to the weather: Lessons from U.S. history,” *The Journal of Economic History*, Vol. 77, No. 3, pp. 756–795.

Burke, Marshall, Solomon M. Hsiang, and Edward Miguel (2015) “Global non-linear effect of temperature on economic production,” *Nature*, Vol. 527, pp. 235–239.

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

Chen, Yuyu, Avraham Ebenstein, Michael Greenstone, and Hongbin Li (2013) “Evidence on the impact of sustained exposure to air pollution on life expectancy from China’s Huai River policy,” *Proceedings of the National Academy of Sciences*, Vol. 110, No. 32, pp. 12936–12941.

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

- Colacito, Riccardo, Bridget Hoffmann, and Toan Phan (2019) “Temperature and growth: A panel analysis of the United States,” *Journal of Money, Credit and Banking*, Vol. 51, No. 2-3, pp. 313–368.
- Costinot, Arnaud, Dave Donaldson, and Cory Smith (2016) “Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world,” *Journal of Political Economy*, Vol. 124, No. 1, pp. 205–248.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (2012) “Temperature shocks and economic growth: Evidence from the last half century,” *American Economic Journal: Macroeconomics*, 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.
- Deschenes, Olivier (2014) “Temperature, human health, and adaptation: A review of the empirical literature,” *Energy Economics*, Vol. 46, pp. 606–619.
- 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.
- (2011) “Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US,” *American Economic Journal: Applied Economics*, Vol. 3, No. 4, pp. 152–185.
- (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.
- Deschênes, Olivier and Enrico Moretti (2009) “Extreme weather events, mortality, and migration,” *The Review of Economics and Statistics*, Vol. 91, No. 4, pp. 659–681.
- Dingel, Jonathan I., Kyle C. Meng, and Solomon M. Hsiang (2019) “Spatial correlation, trade, and inequality: Evidence from the global climate,” Working Paper 25447, National Bureau of Economic Research.
- Duquette, Eric, Nathaniel Higgins, and John Horowitz (2012) “Farmer discount rates: Experimental evidence,” *American Journal of Agricultural Economics*, Vol. 94, No. 2, pp. 451–456.

- 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.
- Gouel, Christophe and David Laborde (2018) “The crucial role of international trade in adaptation to climate change,” Working Paper 25221, National Bureau of Economic Research.
- 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.
- (2014) “Temperature and the allocation of time: Implications for climate change,” *Journal of Labor Economics*, Vol. 32, No. 1, pp. 1–26.
- Guo, Christopher and Christopher Costello (2013) “The value of adaption: Climate change and timberland management,” *Journal of Environmental Economics and Management*, Vol. 65, No. 3, pp. 452–468.
- Hamermesh, Daniel S. (1995) “Comment,” *ILR Review*, Vol. 48, No. 4, pp. 835–838.
- Hamermesh, Daniel S. and Gerard A. Pfann (1996) “Adjustment costs in factor demand,” *Journal of Economic Literature*, Vol. 34, No. 3, pp. 1264–1292.
- Heal, Geoffrey (1976) “The relationship between price and extraction cost for a resource with a backstop technology,” *The Bell Journal of Economics*, Vol. 7, No. 2, pp. 371–378.
- Heal, Geoffrey and Jisung Park (2013) “Feeling the heat: Temperature, physiology and the wealth of nations,” Working Paper 19725, National Bureau of Economic Research.
- (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 (2019) “Beyond random assignment: Credible inference of causal effects in dynamic economies,” *Journal of Finance*, Vol. forthcoming.
- Hsiang, Solomon M. (2016) “Climate econometrics,” *Annual Review of Resource Economics*, Vol. 8, pp. 43–75.
- Huang, Hsing-Hsiang and Michael R. Moore (2019) “Farming under weather risk: Adaptation, moral hazard, and selection on moral hazard,” in Wolfram Schlenker ed. *Understanding Productivity Growth in Agriculture*: University of Chicago Press, pp. 77–124.
- 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.
- 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.
- Kala, Namrata (2017) “Learning, adaptation, and climate uncertainty: Evidence from Indian agriculture,” working paper.
- Keane, Michael P. and Kenneth I. Wolpin (2002) “Estimating welfare effects consistent with forward-looking behavior. Part I: Lessons from a simulation exercise,” *The Journal of Human Resources*, Vol. 37, No. 3, pp. 570–599.
- Kelly, David L., Charles D. Kolstad, and Glenn T. Mitchell (2005) “Adjustment costs from environmental change,” *Journal of Environmental Economics and Management*, Vol. 50, No. 3, pp. 468–495.
- 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.
- Klemm, Toni and Renee A. McPherson (2017) “The development of seasonal climate forecasting for agricultural producers,” *Agricultural and Forest Meteorology*, Vol. 232, pp. 384–399.
- Lemoine, Derek (2017) “Expect above average temperatures: Identifying the economic impacts of climate change,” Working Paper 23549, National Bureau of Economic Research.

- 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.
- Massetti, Emanuele and Robert Mendelsohn (2018) “Measuring climate adaptation: Methods and evidence,” *Review of Environmental Economics and Policy*, Vol. 12, No. 2, pp. 324–341.
- Mendelsohn, Robert (2000) “Efficient adaptation to climate change,” *Climatic Change*, Vol. 45, No. 3-4, pp. 583–600.
- (2019) “An examination of recent revealed preference valuation methods and results,” *Review of Environmental Economics and Policy*, Vol. 13, No. 2, pp. 267–282.
- Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw (1994) “The impact of global warming on agriculture: A Ricardian analysis,” *The American Economic Review*, Vol. 84, No. 4, pp. 753–771.
- Mérel, Pierre and Matthew Gammans (2019) “Climate econometrics: Can the panel approach account for long-run adaptation?” working paper, University of California Davis.
- Nickell, Stephen (1981) “Biases in dynamic models with fixed effects,” *Econometrica*, Vol. 49, No. 6, pp. 1417–1426.
- Nordhaus, William (2013) “Integrated economic and climate modeling,” in Peter B. Dixon and Dale W. Jorgenson eds. *Handbook of Computable General Equilibrium Modeling*, Vol. 1: North Holland, Elsevier B.V. pp. 1069–1131.
- 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.
- Olmstead, Alan L. and Paul W. Rhode (2008) *Creating Abundance: Biological Innovation and American Agricultural Development*, New York, NY: Cambridge University Press.
- (2011) “Responding to climatic challenges: Lessons from U.S. agricultural development,” in Gary D. Libecap and Richard H. Steckel eds. *The Economics of Climate Change: Adaptations Past and Present*, pp. 169–194.
- 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.
- Roberts, Michael J. and Wolfram Schlenker (2011) “The evolution of heat tolerance of corn: Implications for climate change,” in Gary D. Libecap and Richard H. Steckel eds. *The Economics of Climate Change: Adaptations Past and Present*, pp. 225–251.
- Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher (2005) “Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach,” *American Economic Review*, Vol. 95, No. 1, pp. 395–406.
- Schlenker, Wolfram and Michael J. Roberts (2009) “Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change,” *Proceedings of the National Academy of Sciences*, Vol. 106, No. 37, pp. 15594–15598.
- Severen, Christopher, Christopher Costello, and Olivier Deschne (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 (2017) “Expectations and adaptation to environmental risks,” working paper.
- Sohngen, Brent and Robert Mendelsohn (1998) “Valuing the impact of large-scale ecological change in a market: The effect of climate change on U.S. timber,” *The American Economic Review*, Vol. 88, No. 4, pp. 686–710.
- Sorkin, Isaac (2015) “Are there long-run effects of the minimum wage?” *Review of Economic Dynamics*, Vol. 18, No. 2, pp. 306–333.
- Takle, Eugene S., Christopher J. Anderson, Jeffrey Andresen, James Angel, Roger W. Elmore, Benjamin M. Gramig, Patrick Guinan, Steven Hilberg, Doug Kluck, Raymond Massey, Dev Niyogi, Jeanne M. Schneider, Martha D. Shulski, Dennis Todey, and Melissa Widhalm (2013) “Climate forecasts for corn producer decision making,” *Earth Interactions*, Vol. 18, No. 5, pp. 1–8.
- Taraz, Vis (2017) “Adaptation to climate change: historical evidence from the Indian monsoon,” *Environment and Development Economics*, Vol. 22, No. 5, pp. 517–545.
- Weyl, E. Glen (2019) “Price theory,” *Journal of Economic Literature*, Vol. 57, No. 2, pp. 329–384.
- Zhang, Peng, Olivier Deschenes, Kyle Meng, and Junjie Zhang (2018) “Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants,” *Journal of Environmental Economics and Management*, Vol. 88, pp. 1–17.