NBER WORKING PAPER SERIES

EXPECT ABOVE AVERAGE TEMPERATURES: IDENTIFYING THE ECONOMIC IMPACTS OF CLIMATE CHANGE

Derek Lemoine

Working Paper 23549
http://www.nber.org/papers/w23549

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2017

Expect Above Average Temperatures: Identifying the Economic Impacts of Climate Change
Derek Lemoine
NBER Working Paper No. 23549
June 2017
JEL No. D84,H43,Q12,Q51,Q54

ABSTRACT

A rapidly growing empirical literature seeks to estimate the costs of future climate change from time series variation in weather. I formally analyze the consequences of a change in climate for economic outcomes. I show that those consequences are driven by changes in the distribution of realized weather and by expectations channels that capture how anticipated changes in the distribution of weather affect current and past investments. Studies that rely on time series variation in weather omit the expectations channels. Quantifying the expectations channels requires estimating how forecasts affect outcome variables and simulating how climate change would alter forecasts.

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

Climate is what you expect; weather is what you get.[1]

1 Introduction

Economic analyses of climate change policies require estimates of the costs imposed by climate change. Yet we know remarkably little about these costs, leading some economists to question the value of conventional modeling (e.g., Ackerman et al., 2009; Pindyck, 2013). The costs of climate change are hard to pin down before having lived through the climate change experiment. However, we do have access to a rich source of variation in climate variables: we live through changes in the weather on a daily basis. And since climate is just the distribution of weather, many have wondered whether we can substitute this rich source of variation for the missing time series variation in climate.

Pursuing this agenda, a rapidly growing literature has begun to estimate the costs of future climate change by using time series variation in weather.[2] This literature has sought to identify the effects of weather on outcomes such as gross domestic product (Dell et al., 2012), agricultural profits (Deschênes and Greenstone, 2007), crop yields (Schlenker and Roberts, 2009), productivity (Heal and Park, 2013), health (Deschenes, 2014), mortality (Barreca et al., 2016), crime (Ranson, 2014), energy use (Auffhammer and Aroonruengsawat, 2011; Deschênes and Greenstone, 2011), income (Deryugina and Hsiang, 2014), emotions (Baylis, 2017), and more. The typical study first estimates the causal effects of weather on the dependent variable of interest and then uses physical climate models' projections to simulate how the dependent variable would be affected by future climate change. Yet for all the advances this literature has made in connecting weather to various outcomes, the motivating link between weather and climate has lacked theoretical underpinning. Climate is clearly not just weather, but it is indeed just the long-run distribution of weather. What can we learn from the weather about the effects of a change in climate?

I develop a formal model of decision-making under climate change that can guide empirical research. Each period's weather is drawn from a distribution that depends on the climate. Agents' time t payoffs depend on the time t weather realization, their chosen time t controls, and their chosen controls prior to time t. For instance, a farm's time t profit may depend on time t temperature and rainfall, on time t irrigation choices, and on time $t-1$ crop choices. Agents are forward-looking, so their decisions can depend on their beliefs about future weather. Agents observe the current weather before selecting their controls and have access to a forecast of the next period's weather, which draws on knowledge of the current period's weather and of the climate. We are interested in the average effect of a change in

[1] Common variant of Andrew John Herbertson (1901), *Outlines of Physiography*

[2] See Auffhammer and Mansur (2014), Dell et al. (2014), Deschenes (2014), Carleton and Hsiang (2016), and Heal and Park (2016) for surveys. An older literature relied on cross-sectional variation (e.g., Mendelsohn et al., 1994; Schlenker et al., 2005; Nordhaus, 2006), but as discussed in the surveys, cross-sectional approaches have fallen out of favor due to concerns about omitted variables bias.

the climate on agents' intertemporal value, flow payoffs, and optimal choices of controls. In the previous example, a farmer chooses her crop varieties after observing a signal of future temperature and rainfall, and she chooses her quantity of irrigation after observing whether a heat wave or drought has in fact come. We are interested in the average effect of a change in climate on the present discounted value of future profits (as capitalized in land values), on annual profit, and on controls such as irrigation.

I show that a change in climate affects dependent variables of interest through direct weather channels and through expectations channels. The direct weather channels reweight the dependent variable of interest for the new distribution of the weather. Calculating these channels requires estimating how the dependent variable changes with the weather and how the distribution of the weather changes with the climate. This exercise matches the approach followed by the empirical literature to date.[3] The expectations channels account for how a change in the climate alters agents' forecasts of later weather and thereby alters durable investments, such as in crop varietals, levees, or air conditioning. Expectations matter for time t dependent variables in two ways. First, altered expectations of time t weather can affect durable investments in previous periods. Second, altered time t expectations of later weather can affect durable investments at time t. Both of these expectations channels allow for adaptation in advance of a weather shock actually occurring. I show that both expectations channels vanish if agents' actions do not depend on the weather or do not have durable consequences. In these cases, a change in climate does indeed reduce to a change in weather. However, these assumptions are unlikely to apply in general, so empirical work must validate them on a case by case basis.

Estimating the economic consequences of climate change poses both econometric and climate modeling challenges. The econometric challenge of estimating the effect of the weather on dependent variables of interest has been well appreciated, as has the climate modeling challenge of simulating changes in the distribution of the weather. I describe a new pair of challenges. I show that estimating the economic consequences of climate change also requires econometrically estimating how dependent variables of interest change with *forecasts* of the weather and interpreting climate models' output in terms of changes in these forecasts.[4] Applied econometricians have two estimation tasks, and climate modelers must distinguish the forecastable and unforecastable components of future weather.[5]

[3]In practice, the literature has often focused on changes in summary statistics of the weather, such as average temperature.

[4]Some recent work has studied the effects of forecasts (e.g., Rosenzweig and Udry, 2013, 2014; Shrader, 2017). There is also a long literature, primarily in agricultural economics, that seeks to value forecasts and evaluate their usefulness. See Hill and Mjelde (2002), Meza et al. (2008), and Katz and Lazo (2011) for recent surveys. This literature tends to adopt simulation-based approaches (e.g., Solow et al., 1998; Mjelde et al., 2000) rather than the econometric approaches of interest in the present paper.

[5]Distinguishing the forecastable and unforecastable components of future weather is also a task for economists, as it requires estimating how agents form expectations of future weather. Reduced-form approaches have dominated the empirical climate economics literature, but this observation could motivate

I also show that the existence of forecasts can complicate the standard approach to econometrically estimating even the direct effect of the weather. Weather is commonly taken to be exogenous to economic decision-making, in which case time series variation in weather identifies the causal effect of a change in the weather. However, I show that time series estimates are vulnerable to a previously unappreciated omitted variables bias. Time t outcomes often depend on time $t-1$ choices, many of which appear in the error term of standard time series (or panel) regressions. These time $t-1$ choices in turn depend on time $t-1$ weather and on time $t-1$ forecasts of time t weather. Therefore, time series estimates of the causal effect of time t weather on time t dependent variables are biased in the common case where weather is serially correlated and/or partially forecastable.[6] Future work should seek instruments that isolate truly surprising variation in weather, unrelated to either past weather or past forecasts.

As an example, let the variable of interest be time t agricultural profits, which depend on time t temperature and on time $t-1$ choices of crop varieties. The standard approach would regress time t profits on time t temperature, under the assumption that weather, being chosen by nature rather than man, is exogenous to any other factors that might affect profits. However, time $t-1$ crop choices depend on time $t-1$ beliefs about time t temperature. Those beliefs may be influenced by observations of temperature at time $t-1$ and by forecasts released at time $t-1$. Time $t-1$ crop choices clearly affect time t agricultural profits and thus are included in the standard regression's error term. But if temperature is serially correlated between the two periods or if the time $t-1$ forecasts are at all skillful, then these time $t-1$ crop choices are correlated with time t temperature. Past choices can thus act as omitted variables in the standard weather regression. In this example, standard methods may fail to account for unobserved dimensions of crop choices when estimating the effect of temperature on profits, and we already discussed how standard methods fail to recognize expectations-based changes in crop choices when extrapolating the effects of weather shocks to a change in the climate. The first failure potentially works to understate the consequences of a truly surprising weather shock and thus to understate the direct weather component of climate change, but the second failure potentially works to overstate the total costs of climate change. The net bias in standard estimates of the cost of climate change is unclear.

Previous literature has defended the reduction of climate change to time series variation in weather in two ways. First, some authors appeal to the envelope theorem. As presented in Hsiang (2016), the argument is that (1) a change in climate can differ from a change

future use of structural approaches.

[6]Previous work has indeed shown that forecasts matter. Lave (1963) illustrates the value of rain forecasts to raisin growers, and Wood et al. (2014) find that developing-country farmers with better access to weather information make more changes in their farming practices. Roll (1984) and Shrader (2017) both take care to consider weather surprises relative to forecasts. Severen et al. (2016) show that cross-sectional approaches to estimating the dependence of land values on climate have been biased by ignoring priced-in expectations of future climate change. Neidell (2009) demonstrates the importance of accounting for forecasts when estimating the health impacts of air pollution.

in weather only because of differences in beliefs; (2) beliefs can matter only through the choice of control; (3) marginally changing a control must have no effect on payoffs around an optimum; (4) therefore beliefs do not matter for payoffs; (5) therefore a change in climate has identical effects as a change in weather. I highlight two challenges to this argument. First, I show that this argument ignores how expectations of climate change affect past controls. Past controls (i.e., past investments) are taken as given by a time t decision-maker. Step (3) in the argument arises through optimization, but a time t decision-maker can optimize only time t controls. Time t dependent variables can indeed respond to marginal changes in time $t-1$ controls. Second, the envelope theorem is potentially relevant only when dependent variables are objectives, such as streams of utility or profits. Every other dependent variable can vary with a marginal change in even a time t control. In practice, the envelope theorem applies when the dependent variable is either land values or stock prices, as these are the expectation of a stream of profits, and the envelope theorem may apply to annual profits if a particular setting has weak intertemporal linkages. However, as described above, the literature has studied many more dependent variables, including gross domestic product, productivity, health, energy use, and crop yields. All of these dependent variables are either themselves controls or are functions of controls and thus all can be directly influenced by beliefs, even around an optimum.

Previous literature has also sought to transform time series variation in weather into variation in climate by using "long differences" (e.g., Dell et al., 2012; Burke and Emerick, 2016), which many hope better account for long-run adaptation. This approach differs from the simplest time series approach in aggregating outcomes and weather over many timesteps, so that the time index comes to represent, for example, decades rather than years. The hope is that using only longer-run weather variation allows expectations and adaptation to catch up to average weather within a timestep. I here give a structural meaning to long differences. I show that long differences mitigate the omitted variables bias induced by forecasts, but I also show that the usefulness of their results for climate change is unclear. The problem is that long difference estimates entangle the direct effects of weather with the effects of weather on adaptive, durable investments. This entanglement poses a problem because calculating the cost of climate change requires separately analyzing the implications of climate change for forecasts and for realized weather. The structure of expectations embedded in a long difference estimate is almost surely different from the structure of expectations implied by climate change. It is not clear that the information content of a historical sequence of weather shocks approximates the information content of knowing that some type of permanent change in the climate is underway: farmers may learn little from the incremental changes in average weather realized over the last decades but nonetheless may respond strongly once permanent changes in average weather become apparent or become clearly forecasted. Long difference estimates aim to approximate a change in the climate by aggregating the effects of variations in the weather, but when it comes to information sets, a change in the climate could be different even from aggregated variations in the weather.

The next section describes the setting. Section 3 analyzes the effects of climate change on dependent variables of interest and shows which combinations of assumptions reduce a change in climate to the direct weather channels. Section 4 demonstrates underappreciated challenges to causally identifying even the direct weather channels. Section 5 relates the analysis to envelope theorem arguments and long difference approaches. The final section concludes.

2 A Model of Decision-Making Under Climate Change

The minimal model for decision-making under climate change requires three periods: a period of interest, a later period (so that expectations can matter in the period of interest), and an earlier period (so that past expectations can matter for the period of interest). I here develop a general form of such a model.[7]

An agent selects an action a_1 in period 1 and obtains payoffs $\pi_1(a_1, w_1)$ that depend on the action and on the weather w_1, with π_1 strictly concave in a_1. From the perspective of an applied econometrician at time 0, w_1 has probability $p_1(w_1; C)$, where C is a climate index. Period 1 weather is realized before the agent chooses a_1.

In period 2, the agent chooses another action a_2 and receives payoffs $\pi_2(a_1, a_2, w_2)$, which may depend on the period 1 action and on period 2 weather w_2.[8] Let π_2 be strictly concave in a_2 and weakly concave in a_1. The period 2 weather is a random variable from the perspective of the period 1 agent, but that agent has access to a forecast $\theta_1(w_1, C)$ that she uses to update her beliefs about period 2 weather. The forecast may depend on knowledge of the climate and, in light of possible serial correlation in weather, on the period 1 weather outcome. The agent's posterior probability density function for period 2 weather is $p_2(w_2; \theta_1)$. Assume that the period 1 agent correctly extracts information from the forecast, so that she has rational expectations over period 2 weather. The distribution from which w_2 is actually drawn is then described by $p_2(w_2; \theta_1)$. The period 2 weather is realized before the agent chooses her control a_2, and the agent also obtains access to a forecast $\theta_2(w_2, C)$ of period 3 weather w_3 before choosing her control a_2.[9] The agent uses this forecast to assess the distribution of period 3 weather as $p_3(w_3; \theta_2)$.[10]

[7]The most similar model is Kelly et al. (2005). They are interested in the additional costs of having to learn about a change in the climate from an altered sequence of weather as opposed to knowing outright how the climate has changed. They therefore focus on mapping uncertainty about future climate change into the variance of the weather.

[8]I abstract from constraints. One could also model a_1 as affecting constraints on a_2. The results would be qualitatively similar to those we will obtain below.

[9]In many cases, controls must be chosen before the weather is realized. Such cases can be matched to the current setting by interpreting those controls as a_1 rather than a_2.

[10]In general, the distributions of period 2 and period 3 weather could also depend on the climate index C directly, as when an agent's forecast does not include all effects of climate change on the weather distribution. For instance, agents may learn about the effects of climate change over time (e.g., Kelly et al., 2005), which

In period 3, the weather is realized before the agent chooses her control a_3. The period 2 agent had rational expectations, so period 3 weather is in fact distributed as $p_3(w_3; \theta_2)$. The agent receives payoff $\pi_3(a_2, a_3, w_3)$, with π_3 strictly concave in a_3 and weakly concave in a_2.

This setting can capture a variety of stories about climate impacts and adaptation. I give five examples.

1. First, each period could be a year, with the control being the choice of crop to plant. This year's crop choice affects next year's profits when there is a cost to switching crops from year to year. The farmer has access to a drought forecast when making planting decisions.

2. Second, the three periods could occur within a single harvest cycle. Period 1 would then be the spring planting decisions, period 2 would include growing season choices such as irrigation and fertilizer application, and period 3 would represent the harvest. The farmer has access to multiweek or multimonth forecasts when making these decisions.

3. Third, this setting can represent decisions about flood protection. In that case, π_1 would decrease in a_1 so as to capture the costs of, for instance, building levees or raising one's home, and π_2 would increase in a_1 for at least some weather outcomes w_2 so as to capture the benefits of levees or a raised home. The decision-maker has access to forecasts of future rainfall, which determine the expected benefits of flood protection and which may change after observing an unexpectedly large rainfall.

4. Fourth, household investments in air conditioning can provide immediate benefits based on the current weather and can also provide future benefits that depend on future weather. Households may purchase air conditioning based on forecasts of a heat wave in the coming week.

5. Fifth, workers can schedule vacation and tasks around the weather. Weather forecasts of a week or more are now a central feature of daily life. Workers who anticipate, for instance, extreme heat in period 2 can undertake outdoor tasks in period 1 and perhaps even plan to go to the beach or the mountains in period 2. Alternately, office workers in a cold period 1 who anticipate warm weather in period 2 may concentrate their hours and effort into period 1 so as to enjoy the weather in period 2.

In period 3, the agent solves:

$$V_3(a_2, w_3) = \max_{a_3} \pi_3(a_2, a_3, w_3).$$

would prevent their forecasts from capturing the full effect of the climate on the distribution of weather. We here assume that agents and modelers have access to the same information about how climate affects the weather. Extending the setting to allow the modeler to have different information about the effects of climate change (i.e., allowing the modeler and agent to use different weather distributions) would not affect the interaction between the climate and agents' decisions that is of primary interest here.

The first-order condition implicitly defines the optimal period 3 control $a_3^*(a_2, w_3)$:

$$\frac{\partial \pi_3(a_2, a_3^*, w_3)}{\partial a_3} = 0.$$

In period 2, the agent solves:

$$V_2(a_1, w_2) = \max_{a_2} \left\{ \pi_2(a_1, a_2, w_2) + \beta E_2 \left[V_3(a_2, w_3) \right] \right\},$$

where E_t represents expectations at the time t information set (i.e., using θ_t). The first-order condition implicitly defines the optimal period 2 control $a_2^*(a_1, w_2, \theta_2)$:

$$\frac{\partial \pi_2(a_1, a_2^*, w_2)}{\partial a_2} + \beta E_2 \left[\frac{\partial \pi_3(a_2^*, a_3^*, w_3)}{\partial a_2} \right] = 0.$$

And in period 1, the agent solves:

$$V_1(w_1) = \max_{a_1} \left\{ \pi_1(a_1, w_1) + \beta E_1 \left[V_2(a_1, w_2) \right] \right\}.$$

The first-order condition implicitly defines the optimal period 1 control $a_1^*(w_1, \theta_1)$:

$$\frac{\partial \pi_1(a_1^*, w_1)}{\partial a_1} + \beta E_1 \left[\frac{\partial \pi_2(a_1^*, a_2^*, w_2)}{\partial a_1} \right] = 0.$$

We are interested in how period 2 value V_2, payoffs π_2, and controls a_2 change in response to a change in the climate index C, with the applied econometrician's expectations of these changes taken at time 0 (before any weather variables have been realized). We study period 2 outcomes because period 2 is the only period that can enter into agents' expectations while also containing expectations of the future. Economists have tried to identify how climate change will affect all three of these dependent variables: the effect of climate change on land values or stock prices corresponds to changes in V_2, the effect of climate change on profits from growing particular crops corresponds to changes in π_2, and the effect of climate change on decision variables such as hours worked, crime, and air conditioning use correspond to changes in a_2.[11] We now analyze how climate change affects these variables of interest before considering how to econometrically identify the effects of climate change.

[11]The effects of climate change on variables such as gross domestic product, health, and farm yields correspond to changes in a function of a_2, but this function is not typically π_2 or V_2: in standard models, agents (whether households or firms) do not seek to maximize production, health, or yields.

3 The Consequences of Climate Change

Begin by considering how climate change affects expected period 2 value $E_0[V_2]$, which corresponds to dependent variables such as land values. These expectations are taken with respect to time 0 because we, as applied econometricians, seek to identify expected impacts before we know what the weather will be in a given year. We seek the average treatment effect of climate change, where the averaging occurs over weather realizations just as over a population of individuals. This expected value is:

$$E_0\left[V_2(a_1^*(w_1,\theta_1),w_2)\right] = \int\int V_2(a_1^*(w_1,\theta_1),w_2)\,p_2(w_2;\theta_1(w_1,C))\,\mathrm{d}w_2\,p_1(w_1;C)\,\mathrm{d}w_1.$$

If we marginally increase the climate index C, this value changes as:

$$\begin{aligned}\frac{\mathrm{d}E_0\left[V_2\right]}{\mathrm{d}C} =& \int\int \pi_2(a_1^*,a_2^*,w_2)\,\frac{\mathrm{d}p_0(w_1,w_2)}{\mathrm{d}C}\,\mathrm{d}w_2\,\mathrm{d}w_1 \\ &+ \beta\int\int\int \pi_3(a_2^*,a_3^*,w_3)\,\frac{\mathrm{d}p_0(w_1,w_2,w_3)}{\mathrm{d}C}\,\mathrm{d}w_3\,\mathrm{d}w_2\,\mathrm{d}w_1 \\ &+ \int\int \underbrace{\frac{\partial\pi_2(a_1^*,a_2^*,w_2)}{\partial a_1}\frac{\partial a_1^*(w_1,\theta_1)}{\partial\theta_1}}_{\mathrm{d}V_2/\,\mathrm{d}\theta_1}\,\frac{\partial\theta_1(w_1,C)}{\partial C}\,p_0(w_1,w_2)\;\mathrm{d}w_2\,\mathrm{d}w_1, \end{aligned} \tag{1}$$

where we substitute in for $\mathrm{d}V_2/\,\mathrm{d}C$ from the envelope theorem and write $p_0(x,y)$ for the joint distribution of x and y evaluated at time 0. The first two lines are *direct weather channels*. The first line reweights period 2 flow payoffs to reflect changes in the joint distribution of period 1 and period 2 weather. It captures the effect of realized weather on profits, where the period 1 weather realization affects period 2 profits by affecting the period 1 control and where the period 2 weather realization affects period 2 profits both directly and through the period 2 control. The second line reweights period 3 flow payoffs in a similar fashion. It captures the effect of anticipated changes in weather on future profits. Identifying these channels requires identifying how time t flow payoffs change with current and past weather and then using a physical climate model to calculate how climate change may alter the distribution of weather. We will see that in certain special cases one can ignore the effects of past weather on time t payoffs. These special cases are consistent with the type of exercise commonly undertaken in the literature (see Carleton and Hsiang, 2016), though we will discuss in Section 4 why identifying the effect of w_t on π_t is more complicated than commonly recognized.

The third line is a *past expectations channel*. It arises because (and only because) climate change directly alters period 1 beliefs about period 2 weather and thus alters the marginal benefit to period 1 investment. For instance, expectations of different weather due to climate change may drive a period 1 farmer to plant a different crop or invest in an irrigation system, a polity to build levees, or a household to adopt air conditioning.[12] In this interpretation, the past expectations channel captures the benefits of adaptation. One could plausibly construct $\partial\theta_1/\partial C$ from simulations of a climate model, but estimating how either a_1^* or V_2 changes with θ_1 poses a challenge that has been generally overlooked in the empirical climate economics literature to date.

Now consider how a marginal change in the climate index affects expected period 2 payoffs $E_0[\pi_2]$, which corresponds to dependent variables such as profits. We have:

$$\begin{aligned}\frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} =& \int\int \pi_2(a_1^*, a_2^*, w_2)\,\frac{\mathrm{d}p_0(w_1,w_2)}{\mathrm{d}C}\,\mathrm{d}w_2\,\mathrm{d}w_1\\ &+\int\int\left[\frac{\partial\pi_2}{\partial a_1}+\frac{\partial\pi_2}{\partial a_2}\frac{\partial a_2^*(a_1^*,w_2,\theta_2)}{\partial a_1}\right]\frac{\partial a_1^*(w_1,\theta_1)}{\partial\theta_1}\frac{\partial\theta_1(w_1,C)}{\partial C}\,p_0(w_1,w_2)\,\mathrm{d}w_2\,\mathrm{d}w_1\\ &+\int\int\frac{\partial\pi_2}{\partial a_2}\frac{\partial a_2^*(a_1^*,w_2,\theta_2)}{\partial\theta_2}\frac{\partial\theta_2(w_2,C)}{\partial C}\,p_0(w_1,w_2)\,\mathrm{d}w_2\,\mathrm{d}w_1,\end{aligned} \tag{2}$$

where we suppress arguments for π_2 in the last two lines. The first line is a direct weather channel, as in the first line of equation (1). The second line is a past expectations channel, analogous to the third line in equation (1) but now with a new term that accounts for how an expectations-induced change in a_1^* affects π_2 by changing a_2^*. The third line is new. It is a *current expectations channel*. It reflects how the period 2 control depends on expectations of period 3 weather. Climate change affects period 2 payoffs π_2 via period 1 expectations of period 2 weather (second line) and also via period 2 expectations of period 3 weather (third line).[13] One may interpret the past expectations channel as capturing the benefits of adaptation and the current expectations channel as capturing the costs of adaptation.

Finally, consider how a marginal change in the climate index affects the expected period 2 control $E_0[a_2^*]$, which corresponds to dependent variables such as the quantity of irrigation.

[12]In contrast, durable investment decisions mattered in the first line only insofar as climate change altered period 1 weather. For instance, a warm winter may change the timing of planting and thus payoffs during the summer, but that decision does not depend on beliefs about climate change per se. It would arise even for a farmer who knew nothing about climate change and just happened to observe a warm winter with early blooms on the plants. Or a country may build levees or adopt air conditioning because it experienced flooding or a heat wave in period 1, not because beliefs about climate change led it to expect flooding or a heat wave in period 2. Even a nonbeliever in climate change bears the effects reported in the first line, whereas only agents with expectations of climate change bear the effects in the third line. The effects of period 3 weather on period 2 value in the second line may arise even for a nonbeliever if, for instance, the value of land or a firm is determined in a market where other actors do believe in climate change.

[13]Anticipating Section 5, the changes in the period 2 control a_2^* seen on the second and third lines were missing from equation (1) because there the envelope theorem (i.e., period 2 optimization) ensured that we did not need to consider how climate change affects the period 2 choice of control. The envelope theorem applies only to total value V_2, not to flow payoffs π_2. And it never allows us to ignore the effect of climate change on the earlier, period 1 choice of control.

We have:

$$\begin{aligned}\frac{\mathrm{d}E_0\left[a_2^*\right]}{\mathrm{d}C} =& \int\int a_2^*(a_1^*, w_2, \theta_2)\frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C}\,\mathrm{d}w_2\,\mathrm{d}w_1\\ &+\int\int \frac{\partial a_2^*(a_1^*, w_2, \theta_2)}{\partial a_1}\frac{\partial a_1^*(w_1, \theta_1)}{\partial \theta_1}\frac{\partial \theta_1(w_1, C)}{\partial C}\,p_0(w_1, w_2)\;\mathrm{d}w_2\,\mathrm{d}w_1\\ &+\int\int \frac{\partial a_2^*(a_1^*, w_2, \theta_2)}{\partial \theta_2}\frac{\partial \theta_2(w_2, C)}{\partial C}\,p_0(w_1, w_2)\;\mathrm{d}w_2\,\mathrm{d}w_1. \end{aligned} \tag{3}$$

The first line is a direct weather channel, as seen in the first line of equation (1) and also the first line of equation (2). The second line is a past expectations channel, as seen in the third line of equation (1) and also the second line of equation (2). The final line is a current expectations channel, as seen in the third line of equation (2). We need to consider how climate change affects period 2 controls by altering past investments that relied on forecasts of period 2 weather and also by altering period 2 forecasts of future weather.

Thus far, we have seen that estimating the consequences of climate change requires estimating how a time t dependent variable changes with realized time t weather, with past forecasts of time t weather, and, for many dependent variables of interest, with time t forecasts of future weather. We now consider two special cases: when agents cannot mitigate weather shocks, and when agents' decisions do not have long-term consequences. In each of these special cases, we will also explore the implications of the following assumption:

Assumption 1. $\partial^2\theta_t/\partial w_t \partial C = 0$ *for* $t = 1, 2$.

This assumption says that the effect of climate change on time t forecasts (and thus on time $t+1$ weather) is independent of the time t weather realization. It yields the following lemma:

Lemma 1. *Let Assumption 1 hold. Then*

$$\int \frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C}\,\mathrm{d}w_1 = \frac{\partial p_2(w_2;\theta_1)}{\partial \theta_1}\frac{\partial \theta_1(w_1, C)}{\partial C}$$

and

$$\int\int \frac{\mathrm{d}p_0(w_1, w_2, w_3)}{\mathrm{d}C}\,\mathrm{d}w_2\,\mathrm{d}w_1 = \frac{\partial p_3(w_3;\theta_2)}{\partial \theta_2}\frac{\partial \theta_2(w_2, C)}{\partial C}.$$

Proof. See appendix. □

We now turn to the special cases.

3.1 Special Case 1: Agents cannot mitigate weather shocks

Begin by considering a case in which weather may matter for payoffs but does not interact with decisions that an agent can take. For instance, temperatures during the growing season may matter for the quality or quantity of a later harvest, but a farmer without access to irrigation may have little ability to mitigate these consequences. Or especially hot, humid days may be harmful to health, but a household may lack access to electricity that could power air conditioning. These cases correspond to the following restrictions:

Assumption 2. *$\partial^2 \pi_t / \partial a_t \partial w_t = 0$ for $t = 1, 2, 3$ and $\partial^2 \pi_t / \partial a_{t-1} \partial w_t = 0$ for $t = 2, 3$.*

In this special case, the optimized controls are independent of the weather.

The following proposition and corollary describe the effect of climate change on each variable of interest:

Proposition 1. *Let Assumption 2 hold. Then:*

$$\begin{aligned}
\frac{\mathrm{d}E_0\left[a_2^*\right]}{\mathrm{d}C} =& 0, \\
\frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} =& \int \pi_2(a_1^*, a_2^*, w_2) \left(\int \frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C} \, \mathrm{d}w_1 \right) \mathrm{d}w_2, \\
\frac{\mathrm{d}E_0\left[V_2\right]}{\mathrm{d}C} =& \frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} + \beta \int \pi_3(a_2^*, a_3^*, w_3) \left(\int\int \frac{\mathrm{d}p_0(w_1, w_2, w_3)}{\mathrm{d}C} \, \mathrm{d}w_1 \, \mathrm{d}w_2 \right) \mathrm{d}w_3.
\end{aligned}$$

Corollary 1. *Let Assumptions 1 and 2 hold. Then:*

$$\begin{aligned}
\frac{\mathrm{d}E_0\left[a_2^*\right]}{\mathrm{d}C} =& 0, \\
\frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} =& \int \pi_2(a_1^*, a_2^*, w_2) \frac{\partial p_2(w_2; \theta_1)}{\partial \theta_1} \frac{\partial \theta_1(w_1, C)}{\partial C} \, \mathrm{d}w_2, \\
\frac{\mathrm{d}E_0\left[V_2\right]}{\mathrm{d}C} =& \frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} + \beta \int \pi_3(a_2^*, a_3^*, w_3) \frac{\partial p_3(w_3; \theta_2)}{\partial \theta_2} \frac{\partial \theta_2(w_2, C)}{\partial C} \, \mathrm{d}w_3.
\end{aligned}$$

Proof. See appendix. □

We see three interesting results. First, Proposition 1 shows that controls should not respond to a change in climate. This result arises because controls do not respond to weather when Assumption 2 holds. The implied independence of controls from weather and climate can be used to test the plausibility of Assumption 2 when estimating effects on dependent variables such as V_2 and π_2. Second, Proposition 1 shows that Assumption 2 eliminates the expectations channels when analyzing π_2 and V_2. If weather does not affect the choice of control, then forecasts of future weather do not matter for payoffs and total value. Third, the corollary shows that including Assumption 1 eliminates the need to consider weather in the earlier period. When weather realizations do not interact with climate change in determining future weather, then we can calculate how climate change alters the distribution of a given period's weather directly, without needing to simulate how it alters the distribution of a longer sequence of weather.

The special case described by Corollary 1 is consistent with standard approaches in the literature: we can estimate the effect of a change in climate by identifying the causal effect of weather on flow payoffs and then simulating payoffs under the new distribution of weather. However, this special case cannot be motivating the literature: this special case implies that controls are independent of climate change, but much of the literature has in fact focused on estimating the effects of climate change on controls or on functions of controls. The next special case is probably closer to the spirit of the literature to date.

3.2 Special Case 2: Agents' decisions do not have long-term consequences

We now consider a special case in which decisions affect only contemporaneous payoffs. This restriction turns the setting into a sequence of static investment choices, as when farmers can costlessly switch between crops after each year or when increasing air conditioning requires turning on an existing system rather than installing a new system. Formally, we impose the following restriction:

Assumption 3. *$\partial \pi_t / \partial a_{t-1} = 0$ for $t = 2, 3$.*

This restriction yields the following results:

Proposition 2. *Let Assumption 3 hold. Then:*

$$
\begin{aligned}
\frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} &= \int \pi_2(a_1^*, a_2^*, w_2) \left(\int \frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C} \, \mathrm{d}w_1 \right) \mathrm{d}w_2, \\
\frac{\mathrm{d}E_0\left[V_2\right]}{\mathrm{d}C} &= \frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} + \beta \int \pi_3(a_2^*, a_3^*, w_3) \left(\int \int \frac{\mathrm{d}p_0(w_1, w_2, w_3)}{\mathrm{d}C} \, \mathrm{d}w_1 \, \mathrm{d}w_2 \right) \mathrm{d}w_3, \\
\frac{\mathrm{d}E_0\left[a_2^*\right]}{\mathrm{d}C} &= \int a_2^*(a_1^*, w_2, \theta_2) \left(\int \frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C} \, \mathrm{d}w_1 \right) \mathrm{d}w_2.
\end{aligned}
$$

Corollary 2. *Let Assumptions 1 and 3 hold. Then:*

$$
\begin{aligned}
\frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} &= \int \pi_2(a_1^*, a_2^*, w_2) \, \frac{\partial p_2(w_2; \theta_1)}{\partial \theta_1} \frac{\partial \theta_1(w_1, C)}{\partial C} \, \mathrm{d}w_2, \\
\frac{\mathrm{d}E_0\left[V_2\right]}{\mathrm{d}C} &= \frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} + \beta \int \pi_3(a_2^*, a_3^*, w_3) \, \frac{\partial p_3(w_3; \theta_2)}{\partial \theta_2} \frac{\partial \theta_2(w_2, C)}{\partial C} \, \mathrm{d}w_3, \\
\frac{\mathrm{d}E_0\left[a_2^*\right]}{\mathrm{d}C} &= \int a_2^*(a_1^*, w_2, \theta_2) \frac{\partial p_2(w_2; \theta_1)}{\partial \theta_1} \frac{\partial \theta_1(w_1, C)}{\partial C} \, \mathrm{d}w_2.
\end{aligned}
$$

Proof. See appendix.

□

These results are similar to the results in Section 3.1, with the important difference being that the optimized control a_2^* can now change with the climate. Once we restrict attention to static investments, we lose all of the expectations channels. And with the addition of Assumption 1, we are once again in a special case where, corresponding to much of the recent empirical climate economics literature (see Carleton and Hsiang, 2016), we need only estimate how a dependent variable changes with contemporary weather and then reweight outcomes for the new distribution of weather. The combination of Assumptions 1 and 3 can indeed motivate standard approaches. However, Assumption 3 may be overly restrictive in many cases of interest. Future empirical work should highlight the extent to which a given environment involves dynamic decision-making and test the restrictions imposed by Assumption 3.

4 Identifying the Effects of Weather and Forecasts

We have seen that determining the economic impacts of climate change involves (1) a climate modeling challenge and (2) an econometric challenge. The climate modeling challenge is (1a) to project how climate change alters the distribution of weather and (1b) to project how climate change alters forecasts of the weather. The economics and climate science literatures have devoted substantial attention to (1a) (e.g., Auffhammer et al., 2013; Kirtman et al., 2013; Burke et al., 2014; Melillo et al., 2014; Lemoine and Kapnick, 2016) but not much attention to (1b), except insofar as climate models are themselves forecasts.[14] The econometric challenge is (2a) to identify how the dependent variable of interest changes with the weather and (2b) to identify how the dependent variable of interest changes with forecasts of the weather. The economics literature has devoted substantial attention to (2a), as the many recent reviews attest (Auffhammer and Mansur, 2014; Dell et al., 2014; Deschenes, 2014; Carleton and Hsiang, 2016; Heal and Park, 2016; Hsiang, 2016), but little attention to (2b).[15] However, I will here argue that these estimates of (2a) are not as clearly identified as commonly believed. Intuitively, time t dependent variables often respond to time $t-1$ decisions. Time $t-1$ decisions in turn often depend on either time $t-1$ weather or on time $t-1$ forecasts of time t weather. Because time t weather is typically correlated with time

[14] As an exception, Lemoine and Kapnick (2016) allow forecasts to evolve with the climate. However, they explore only simple forecasting rules rather than modeling forecasts directly, as they are interested in projecting the costs of changing the variance of the climate. Climate modeling studies have found that the occurrence of climate change makes forecasts more skillful, insofar as accounting for increases in greenhouse gases (and the resulting warming trend) produces better results than does using the historical distribution of weather and climate (e.g., Smith et al., 2007; Jia et al., 2014; Yang et al., 2015).

[15] As an exception, Shrader (2017) studies (2b) with fishery revenue as the dependent variable and El Niño as the weather pattern of interest.

time $t-1$ weather and with time $t-1$ forecasts of time t weather, the effect of unobserved time $t-1$ decisions on time t outcomes induces correlation between the time t weather and the error term in the standard time t regression equation. Past weather and information can each induce omitted variables bias in the standard weather regression.

The recent benchmark for estimating the causal impact of weather on a dependent variable y is to estimate an equation of the form:

$$y_{it} = \alpha_i + w_{it}\,\beta + x_{it}\gamma + \nu_t + \epsilon_{it}, \tag{4}$$

where x_{it} is a vector of covariates and β is the coefficient of interest. See, for instance, Hsiang (2016). In this panel regression with time and unit fixed effects, the causal effect of weather w on outcomes y is identified by idiosyncratic weather shocks with respect to the average weather experienced by unit i. In common usage, the coefficient β is identified if and only if $Cov(w_{it}, \epsilon_{it}) = 0$, so that changes in the weather are independent of other changes that could affect the outcome variable. Since weather is often taken to be the ultimate exogenous variable, generated by nature as from random dice rolls, the assumption that $Cov(w_{it}, \epsilon_{it}) = 0$ is often taken to be a safe one.[16]

Now consider the connection to our general theoretical framework. The outcome variable y_{it} could be $V_2(a_1^*, w_2)$, $\pi_2(a_1^*, a_2^*(a_1^*, w_2, \theta_2(w_2, C)), w_2)$, $a_2^*(a_1^*, w_2, \theta_2(w_2, C))$, or a function of these. In all cases, we can write the outcome of interest Y_2 as a function of a_1^* and w_2: $Y_2(a_1^*, w_2)$, with C a parameter. A first-order Taylor series expansion of $Y_2(a_1^*, w_2)$ around some point $(\bar{a}, \bar{w})$ yields:

$$\begin{aligned}
Y_2(a_1^*, w_2) &\approx Y_2(\bar{a}, \bar{w}) + \left.\frac{\partial Y_2(a_1^*, w_2)}{\partial a_1}\right|_{(\bar{a},\bar{w})} (a_1^* - \bar{a}) + \left.\frac{\partial Y_2(a_1^*, w_2)}{\partial w_2}\right|_{(\bar{a},\bar{w})} (w_2 - \bar{w}) \\
&= \left[Y_2(\bar{a}, \bar{w}) - \left.\frac{\partial Y_2(a_1^*, w_2)}{\partial a_1}\right|_{(\bar{a},\bar{w})} \bar{a} - \left.\frac{\partial Y_2(a_1^*, w_2)}{\partial w_2}\right|_{(\bar{a},\bar{w})} \bar{w} \right] \\
&\quad + \left.\frac{\partial Y_2(a_1^*, w_2)}{\partial a_1}\right|_{(\bar{a},\bar{w})} a_1^* + \left.\frac{\partial Y_2(a_1^*, w_2)}{\partial w_2}\right|_{(\bar{a},\bar{w})} w_2.
\end{aligned}$$

Matching this to the regression equation (4), we seek $\beta = \left.\frac{\partial Y_2(a_1^*, w_2)}{\partial w_2}\right|_{(\bar{a},\bar{w})}$. The term in brackets on the second line is just a constant that will be absorbed into the fixed effect α_i.

[16] Note that it may or may not be a problem if changes in the weather cause changes in other variables that in turn cause changes in the dependent variable of interest. For instance, recent empirical literature has argued that weather affects labor productivity (Heal and Park, 2013), income (Deryugina and Hsiang, 2014), and economic growth (Dell et al., 2012), which are often omitted variables in regressions of, for instance, the effect of weather on crime (Ranson, 2014). The extent to which this omission poses a problem depends on whether a researcher is interested in the "direct" effect of weather on y_{it} or in the total effect, including indirect effects that arise through effects on omitted variables. In the latter case, one needs to beware of double-counting when tallying up estimated impacts across studies.

The terms on the third line can vary over time and over agents and thus would not be picked up by time or unit fixed effects. The potential problem arises from the dependence of Y_2 on a_1^*. In principle, a_1^* could be included in the vector x_{it}, but the weather regression literature rarely (if ever) mentions including previous decisions in the observed covariates, and even if one were to take that approach, it would be difficult to prove that all relevant controls a_1 are observed and included in x_{it}. Thus, it is likely that $\frac{\partial Y_2(a_1^*, w_2)}{\partial a_1}\Big|_{(\bar{a},\bar{w})} a_1^*$ ends up in ϵ_{it}.

Now consider whether weather shocks really are orthogonal to the error. Assume that

$$\epsilon_{it} = \frac{\partial Y_2(a_1^*, w_2)}{\partial a_1}\Big|_{(\bar{a},\bar{w})} a_1^* + z_{it}, \tag{5}$$

with $Cov(z_{it}, \epsilon_{it}) = 0$.[17] In other words, assume that β is identified unless a_1^* poses some particular problem. This is the best possible case. Using a first-order Taylor series expansion of $a_1^*(w_1, \theta_1)$ around some point $(\bar{w}_1, \bar{\theta}_1)$, we have:

$$\begin{aligned} Cov(w_{it}, \epsilon_{it}) &= \frac{\partial Y_2(a_1^*, w_2)}{\partial a_1}\Big|_{(\bar{a},\bar{w})} Cov(w_2, a_1^*(w_1, \theta_1)) \\ &\approx \frac{\partial Y_2(a_1^*, w_2)}{\partial a_1}\Big|_{(\bar{a},\bar{w})} Cov\left(w_2, \frac{\partial a_1^*(w_1, \theta_1)}{\partial w_1}\Big|_{(\bar{w}_1,\bar{\theta}_1)} w_1 + \frac{\partial a_1^*(w_1, \theta_1)}{\partial \theta_1}\Big|_{(\bar{w}_1,\bar{\theta}_1)} \theta_1\right) \\ &= \frac{\partial Y_2(a_1^*, w_2)}{\partial a_1}\Big|_{(\bar{a},\bar{w})} \frac{\partial a_1^*(w_1, \theta_1)}{\partial w_1}\Big|_{(\bar{w}_1,\bar{\theta}_1)} Cov\left(w_2, w_1\right) \\ &\quad + \frac{\partial Y_2(a_1^*, w_2)}{\partial a_1}\Big|_{(\bar{a},\bar{w})} \frac{\partial a_1^*(w_1, \theta_1)}{\partial \theta_1}\Big|_{(\bar{w}_1,\bar{\theta}_1)} Cov\left(w_2, \theta_1\right). \end{aligned}$$

Weather is often serially correlated, so that $Cov(w_2, w_1) > 0$. And if forecasts have any informational content, then $Cov(w_2, \theta_1) \neq 0$. In these common cases, we need at least one of the following two conditions to hold in order for equation (4) to properly identify β:

1. $\partial Y_2(a_1^*, w_2)/\partial a_1 = 0$,
2. $\partial a_1^*(w_1, \theta_1)/\partial w_1 = 0$ and $\partial a_1^*(w_1, \theta_1)/\partial \theta_1 = 0$.

For proper identification, we need past controls not to matter for the dependent variable of interest, or we need past choices of controls to be independent of past weather and independent of past forecasts. When past controls are irrelevant for time t outcomes, then past information and weather are irrelevant for time t outcomes and do not affect the time

[17] Note that the assumption that $Cov(z_{it}, \epsilon_{it}) = 0$ is violated if (in a more general analysis) the vector w_{it} does not capture all of the relevant weather variables. The presence of correlated weather variables in ϵ_{it} would bias the estimated β. For instance, heat and humidity may be correlated and may both affect outcome variables (e.g., Barreca, 2012), yet many regressions include only temperature in w_{it}.

t regression error.[18] When past choices of controls are independent of past weather and independent of past forecasts, then past controls may appear in the time t regression error but will be uncorrelated with time t weather. In all other cases, the standard weather regression equation (4) is vulnerable to omitted variables bias.[19] As an example, imagine that agents undertake durable adaptation investments in response to extreme weather and/or to forecasts of future weather, as when ongoing droughts lead to investments in irrigation systems. These investments would directly affect later outcomes while also being correlated with later weather, as when past investments in irrigation systems affect current yields while being correlated with current drought conditions through past drought conditions and past drought forecasts.

The following proposition describes sufficient conditions for equation (4) to properly identify β:

Proposition 3. *Let either Assumption 2 or Assumption 3 hold, and let equation* (5) *hold with* $Cov(z_{it}, \epsilon_{it}) = 0$. *Then the coefficient* β *in the regression equation* (4) *is identified.*

Proof. First, let Assumption 2 hold. The proof of Proposition 1 shows that $\partial a_1^*(w_1, \theta_1)/\partial w_1 = 0$ and $\partial a_1^*(w_1, \theta_1)/\partial \theta_1 = 0$.

Second, let Assumption 3 hold. This assumption directly states that π_2 is independent of a_1^*. Using equation (A-3), Assumption 3 also implies that a_2^* is independent of a_1. Therefore, $\partial Y_2(a_1^*, w_2)/\partial a_1 = 0$ for each possible meaning of Y_2.

□

The regression equation (4) is identified in the same special cases in which the effect of climate on the outcome of interest reduces to the effect of weather on the outcome of interest, without expectations channels. Thus, if changing the climate is economically equivalent to a surprise change in today's weather, whether because agents cannot take decisions that mitigate the consequences of weather shocks or because all decisions are short-run decisions, then the conventional regression equation (4) is identified and also gives us all the information we need to calculate the economic consequences of climate change. However, the regression equation (4) is not identified in general. One may have hoped that the empirical climate economics literature has at least been properly estimating the direct weather channels even

[18] Assuming that past weather cannot directly affect time t outcomes. Any such relation could be accommodated in a more general form of the analysis by letting w_t be a vector that includes past weather realizations.

[19] For instance, Miller (2015) provides evidence that farmers in India select their crops as if they had a signal of the coming season's precipitation. Realized precipitation is thus endogenous in regressions that aim to identify its effect on, for instance, income. Also studying Indian agriculture, Rosenzweig and Udry (2013) show that farmers' investments respond to forecasts (and respond more strongly to more skillful forecasts), and Rosenzweig and Udry (2014) show that forecasts of planting season weather affect migration decisions and thus wages.

when neither Assumption 2 nor Assumption 3 holds, but we now see that the regressions used to identify the consequences of changes in the weather may themselves be biased.

How can we identify β in settings with dynamic and/or weather-dependent investments? First, we could instrument for w_{it}, seeking the unpredictable portion of weather variation. One could imagine using deviations from published forecasts or deviations from retrospective weather models' predictions as the weather outcome of interest. The burden of the argument would rest in establishing that agents did not have access to further sources of information. Second, we could find settings in which the control a_1 is either observed by the applied econometrician or was exogenously fixed independent of weather or of expectations. Quasi-experiments in which a_1 were held fixed for some actors but not for others would reveal the bias present in more naive regressions. The burden of the argument would rest in establishing that the regression's covariates includes all possibly relevant period 1 actions.

How is ignoring forecasts likely to bias standard panel estimates of $\partial Y_2/\partial w_2$? Much of the literature estimates negative impacts of extreme weather on variables of interest. Imagine that a forecast of extreme weather induces a choice of control that mitigates the later negative impact, as when forecasts of heat waves lead households to purchase air conditioning or farmers to supply more water to their crops. In this case, conflating the effects of the forecasts and the weather shock is likely to bias the estimated effect of the weather shock towards zero: protective actions reduce the impact of the weather shock.[20] However, it is unclear how calculations of climate change impacts would be biased by ignoring forecast information. On the one hand, reducing climate change to only the direct weather channels may overstate the costs of climate change by ignoring the potential for adaptation, but on the other hand, estimating the costs of weather shocks from variation that includes a forecasted component may tend to understate the costs of weather shocks and thus to understate the costs of climate change.[21]

Finally, note that estimating the marginal effect of the climate on outcomes of interest also generally requires estimating the marginal effect of forecasts on outcomes of interest. Future empirical analysis that seeks to identify the marginal effects of forecasts will run into similar identification challenges as just discussed above. For instance, consider the following

[20] Alternately, an office worker who applies more effort on cold days so as to take advantage of warmer days to come will bias estimates of the effect of higher temperature on productivity towards a more negative effect.

[21] Recent literature has emphasized that the effects of temperature may be nonlinear, with especially high temperatures causing especially severe damages (e.g., Burke et al., 2015). This result is consistent with the bias story outlined here. Extreme outcomes may be less likely to be correctly forecasted than are more common outcomes, so that the panel regression may come closer to correctly identifying the direct costs of extreme temperatures. However, the same mechanism would imply that the estimated effects of extreme temperatures may be especially uninformative about the costs of climate change, which converts formerly rare extremes into more common, forecastable events.

regression, where we seek $\tilde{\beta}$:

$$y_{it} = \tilde{\alpha}_i + \theta_{i(t-1)}\,\tilde{\beta} + x_{it}\tilde{\gamma} + \tilde{\nu}_t + \tilde{\epsilon}_{it}. \tag{6}$$

It is easy to see that identifying the marginal effect of a better forecast faces a similar challenge as does identifying the marginal effect of weather: past weather w_{t-1} is an input to past forecasts and past forecasts aim to predict current weather w_t, so bias arises whenever current weather can affect current outcomes directly. To identify the $\tilde{\beta}$ in equation (6), we would need to argue that the regression includes every weather channel as an observed covariate in x_{it},[22] or we would need to instrument for θ_{t-1} by using variation in forecasts that does not rely on past weather and does not successfully predict future weather. Farmer's almanacs or long-run hurricane forecasts could be two such sources of exogenous variation. To my knowledge, this type of instrumental variables analysis has not been undertaken to date.

5 Relation to Previous Arguments

We have seen that the effect of a change in climate does not include expectations channels when there are no dynamic linkages between periods and also when weather does not interact with decision-making. Neither of these restrictions is likely to apply to many (perhaps most) cases of interest. Yet a large and growing literature estimates climate impacts from panel variation in weather. This literature has often been subject to the informal complaint that "climate is not weather." We now consider arguments that have been deployed in defense of reducing climate to weather.

5.1 Appeals to the Envelope Theorem

The most forceful and complete response is in Hsiang (2016). He argues that "the total effect of climate can be exactly recovered using $\hat{\beta}_{TS}$ derived from weather variation" (pg. 53). He considers agents who solve the following static problem:

$$Y_t(C) = \max_{b_t} z_t(b_t, w_t(C)), \tag{7}$$

where we add a time subscript to his notation, use w_t for weather in place of his c, and strip away noncritical vector notation. He writes weather as a function of climate C so as to indicate that the weather is drawn from a distribution controlled by C. Totally differentiating, he obtains:

$$\frac{\mathrm{d}Y_t(C)}{\mathrm{d}C} = \frac{\partial z_t(b_t^*(C), w_t(C))}{\partial w_t(C)}\frac{\mathrm{d}w_t(C)}{\mathrm{d}C} + \frac{\partial z_t(b_t^*(C), w_t(C))}{\partial b_t}\frac{\mathrm{d}b_t^*(C)}{\mathrm{d}C}.$$

[22] In his study of the role that El Niño forecasts play in fishery revenue, Shrader (2017) includes realized sea surface temperature as a covariate.

He calls the first term "direct effects" and the second term "belief effects," because the second term depends on how agents adjust their actions as a result of understanding that the climate has changed. The first-order condition for the agent's problem (7) implicitly defines $b_t^*(C)$:

$$\frac{\partial z_t(b_t^*, w_t(C))}{\partial b_t} = 0.$$

Therefore, we have:

$$\frac{\mathrm{d}Y_t(C)}{\mathrm{d}C} = \frac{\partial z_t(b_t^*(C), w_t(C))}{\partial w_t(C)} \frac{\mathrm{d}w_t(C)}{\mathrm{d}C} = \frac{\partial Y_t(C)}{\partial w_t} \frac{\mathrm{d}w_t(C)}{\mathrm{d}C},$$

which is an application of the envelope theorem. The belief effects have vanished because only actions can depend on beliefs and agents maximize their actions. The consequences of any marginal change in the climate can therefore be approximated by estimating how Y_t varies with the weather w_t and then simulating changes in the weather that correspond to a change in the climate. Further, Hsiang (2016) argues that panel regressions of the form in equation (4) can recover $\partial Y_t / \partial w_t$ as the estimated β. By this reasoning, panel regressions tell us not just about the effects of weather on the dependent variable Y_t but also about the effect of climate.

We have seen that, in a dynamic model, beliefs can enter into $\mathrm{d}Y_t / \mathrm{d}C$ through effects on b_{t-1} and through effects on the distribution of z_{t+1}. These channels are missing from the static model considered here. However, there is another problem with the envelope theorem argument, one that is internal to the setting. The envelope theorem argument identifies the y_{it} from the left-hand side of the panel regression (4) with the value function on the left-hand side of the agent's problem (7). Converted to our notation, the maximization problem (7) becomes

$$V_t = \max_{a_t} \pi_t(a_t, w_t(C)).$$

The y_{it} from the regression equation (4) could be identified with V_t, with a_t^*, or with some function thereof.[23] The envelope theorem applies only when we identify y_{it} with V_t.[24] There is no theorem that says that the optimal choice of control (a_t^* in our setting, or b_t^* in Hsiang (2016)) cannot respond to changes in a parameter such as C.

Is it more plausible that y_{it} corresponds to V_t or more plausible that y_{it} corresponds to a_t^*? There are a limited set of options for objective functions in neoclassical settings: individuals and households maximize utility, and firms maximize profits. The applied econometrician does not observe utility. All observed individual-level dependent variables must be controls or functions of controls: settings that estimate the effects of climate on outcomes such

[23] In our dynamic setting, there is a difference between V_t and maximized π_t, so we there saw yet another possible definition for y_{it}.

[24] As Hsiang (2016, 57) recognizes: "Existing papers leveraging weather variation do not explicitly check the assumptions critical to this result: that Y is the solution to a (constrained) maximization..."

as labor productivity (Heal and Park, 2013), health (Deschenes, 2014), mortality (Barreca et al., 2016), crime (Ranson, 2014), energy use (Auffhammer and Aroonruengsawat, 2011; Deschênes and Greenstone, 2011), income (Deryugina and Hsiang, 2014), gross domestic product (Dell et al., 2012), and emotions (Baylis, 2017) cannot rely on envelope theorem arguments.[25] With respect to firms, the applied econometrician can observe profits. Settings such as Deschênes and Greenstone (2007) that focus on profits can appeal to an envelope theorem argument.[26] However, settings that estimate the effects of climate change on yields (Schlenker and Roberts, 2009), input choices (Zhang et al., 2016), and production (Cachon et al., 2012) cannot rely on envelope theorem arguments.

Our dynamic setting in fact shows precisely how far one can get with an envelope theorem argument. In Section 3, we saw that:

1. $\mathrm{d}E_0[\pi_2]/\,\mathrm{d}C$ and $\mathrm{d}E_0[a_2^*]/\,\mathrm{d}C$ each depend in general on $\partial a_2^*/\partial\theta_2$, and
2. $\mathrm{d}E_0[V_2]/\,\mathrm{d}C$, $\mathrm{d}E_0[\pi_2]/\,\mathrm{d}C$, and $\mathrm{d}E_0[a_2^*]/\,\mathrm{d}C$ each depend in general on $\partial a_1^*/\partial\theta_1$,

where a_t^* is the optimized period t control and θ_t is the climate-dependent period t forecast of period $t+1$ weather. Dependence on $\partial a_t^*/\partial\theta_t$ indicates sensitivity to beliefs about future weather. The envelope theorem eliminates $\partial a_2^*/\partial\theta_2$ only from $\mathrm{d}E_0[V_2]/\,\mathrm{d}C$: we can ignore the effect of time t beliefs on the time t control when the dependent variable is a measure of time t intertemporal value (such as land values). However, other dependent variables depend on contemporaneous beliefs, and even intertemporal value can depend on *past* beliefs via past choices of controls. We cannot assume away expectations as a general rule; instead, we must justify that particular settings satisfy restrictions that render expectations irrelevant to the effect of the climate on the dependent variable of interest. We can, at best, abstract from expectations on a case by case basis.

5.2 Use of Long Differences

Sensitive to the criticism that time series variation may not account for expectations-based adaptation, several papers have adopted "long differences" approaches to estimation. As described in Dell et al. (2014) and Burke and Emerick (2016), long differences change the regression equation (4) to

$$y_{id} = \hat{\alpha}_i + w_{id}\,\hat{\beta} + x_{id}\hat{\gamma} + \hat{\nu}_d + \hat{\epsilon}_{id}, \tag{8}$$

where we have replaced the time index t with an index d whose increments correspond to nt units of time. A long difference approach selects n to be a large, positive integer and

[25] For instance, Barreca et al. (2016) demonstrate the importance of air conditioning in mediating the temperature-mortality relationship. This result highlights the potential importance of expectations-driven investments ("adaptation") in distinguishing a change in the climate from a change in the weather.

[26] Indeed, the envelope theorem argument in Hsiang (2016) can be seen as an extension of the line of reasoning begun in Deschênes and Greenstone (2007).

averages y_{it}, w_{it}, and x_{it} over each timestep d. In the prototypical example with only two values of d (e.g., Dell et al., 2012), the regression acts like differencing over the two "long" timesteps. The hope is that the estimated coefficients capture longer-run opportunities for belief formation and adaptation because each timestep now includes longer-run weather variation.

Do long differences accomplish their goal of including expectations channels in their estimates of weather impacts? Consider long differences within our theoretical framework. In order to allow for "differences" of arbitrary "length", extend the setting of Section 2 to an arbitrary number of periods in the natural way. Following Section 4, write $Y_t(a^*_{t-1}, w_t)$ as the reduced-form representation of the outcome of interest, whether that outcome be V_t, π_t, a^*_t, or a function thereof. Consider a long-difference aggregation from t to $T \triangleq nt$:

$$Y_{t\to T}(\mathbf{w}) \triangleq \sum_{s=t}^{T} Y_s(a^*_{s-1}(w_{s-1}, \theta_{s-1}(w_{s-1}, C)), w_s),$$

where $\mathbf{w}$ is the vector of w_s for $s \in \{t-1, t, ..., T\}$. Let $\bar{w}$ be the average weather for the unit of interest. A first-order Taylor series expansion of $Y_{t\to T}$ around $\bar{w}$ yields:

$$\begin{aligned}
Y_{t\to T}(\mathbf{w}) \approx\, & Y_{t\to T}(\bar{w}) \\
& + \sum_{s=t}^{T} \Bigg[\left.\frac{\partial Y_s(a^*_{s-1}, w_s)}{\partial w_s}\right|_{\bar{w}} (w_s - \bar{w}) \\
& \quad + \left.\frac{\partial Y_s(a^*_{s-1}, w_s)}{\partial a_{s-1}}\right|_{\bar{w}} \left(\left.\frac{\partial a^*_{s-1}(w_{s-1}, \theta_{s-1})}{\partial w_{s-1}}\right|_{\bar{w}} + \left.\frac{\partial a^*_{s-1}(w_{s-1}, \theta_{s-1})}{\partial \theta_{s-1}}\right|_{\bar{w}} \left.\frac{\partial \theta_{s-1}}{\partial w_{s-1}}\right|_{\bar{w}} \right) (w_{s-1} - \bar{w}) \Bigg] \\
=\, & \chi_{t\to T} + \left.\frac{\partial Y_T(a^*_{T-1}, w_T)}{\partial w_T}\right|_{\bar{w}} w_T \\
& + \sum_{s=t}^{T-1} \Bigg[\left.\frac{\partial Y_s(a^*_{s-1}, w_s)}{\partial w_s}\right|_{\bar{w}} + \left.\frac{\partial Y_{s+1}(a^*_s, w_{s+1})}{\partial a_s}\right|_{\bar{w}} \left(\left.\frac{\partial a^*_s(w_s, \theta_s)}{\partial w_s}\right|_{\bar{w}} + \left.\frac{\partial a^*_s(w_s, \theta_s)}{\partial \theta_s}\right|_{\bar{w}} \left.\frac{\partial \theta_s}{\partial w_s}\right|_{\bar{w}} \right) \Bigg] w_s \\
& + \left.\frac{\partial Y_t(a^*_{t-1}, w_t)}{\partial a_{t-1}}\right|_{\bar{w}} \left(\left.\frac{\partial a^*_{t-1}(w_{t-1}, \theta_{t-1})}{\partial w_{t-1}}\right|_{\bar{w}} + \left.\frac{\partial a^*_{t-1}(w_{t-1}, \theta_{t-1})}{\partial \theta_{t-1}}\right|_{\bar{w}} \left.\frac{\partial \theta_{t-1}}{\partial w_{t-1}}\right|_{\bar{w}} \right) w_{t-1},
\end{aligned}$$

where $\chi_{t\to T}$ is a constant for given $\bar{w}$. The second-to-last line captures how realized weather w_s affects Y_s (inclusive of the effect of w_s on the control a^*_s) and it captures how realized weather w_s affects Y_{s+1} through expectations-driven investment. The last line reflects how the weather at time $t-1$ affects outcomes at time t by affecting investments at time $t-1$. This term was the source of the bias described in Section 4.

Now assume that the flow payoff π_t and the probabilities $p_t(w_t; \theta_{t-1})$ do not depend directly on time, so that we can drop the time subscript on Y. Write a^* to indicate the steady-state a that would arise from observing $\bar{w}$ forever. We have:

$$\begin{aligned} Y_{t\to T}(\mathbf{w}) \approx{}& \chi_{t\to T} \\ &+ \left[\left.\frac{\partial Y(a^*,w)}{\partial w}\right|_{\bar{w}} + \left.\frac{\partial Y(a^*,w)}{\partial a}\right|_{\bar{w}} \left(\left.\frac{\partial a^*(w,\theta)}{\partial w}\right|_{\bar{w}} + \left.\frac{\partial a^*(w,\theta)}{\partial \theta}\right|_{\bar{w}} \left.\frac{\partial \theta}{\partial w}\right|_{\bar{w}} \right) \right] \sum_{s=t}^{T-1} w_s \\ &+ \left.\frac{\partial Y(a^*,w)}{\partial w}\right|_{\bar{w}} w_T + \left.\frac{\partial Y(a^*,w)}{\partial a}\right|_{\bar{w}} \left(\left.\frac{\partial a^*(w,\theta)}{\partial w}\right|_{\bar{w}} + \left.\frac{\partial a^*(w,\theta)}{\partial \theta}\right|_{\bar{w}} \left.\frac{\partial \theta}{\partial w}\right|_{\bar{w}} \right) w_{t-1}, \end{aligned} \tag{9}$$

noting that we no longer need time subscripts on the θ or the w evaluated at known points. The contribution of the last line becomes small as T becomes large (i.e., as the "difference" becomes "long"), in which case we have:

$$Y_{t\to T}(\mathbf{w}) \approx \chi_{t\to T} + \underbrace{\left[\left.\frac{\partial Y(a^*,w)}{\partial w}\right|_{\bar{w}} + \left.\frac{\partial Y(a^*,w)}{\partial a}\right|_{\bar{w}} \left(\left.\frac{\partial a^*(w,\theta)}{\partial w}\right|_{\bar{w}} + \left.\frac{\partial a^*(w,\theta)}{\partial \theta}\right|_{\bar{w}} \left.\frac{\partial \theta}{\partial w}\right|_{\bar{w}} \right) \right]}_{\Gamma} \sum_{s=t}^{T-1} w_s.$$

The approximation becomes exact as $T \to \infty$, yielding the cross-sectional result in which average weather is all that matters.[27] Helpfully, we no longer have to worry about the bias from ignoring correlation between the error term in (8) and investments at time $t-1$. The coefficient $\hat{\beta}$ should therefore converge to Γ (the term in square brackets), which clearly includes expectations-driven investments (often referred to as "long-run adaptation"). One might therefore hope that we have solved the problem of properly identifying the causal effect of weather at the same time as we have defined an effect that is closer to the experience of changing the climate.

Unfortunately, matters are not so simple, because the effects of forecasts are entangled with the effects of weather inside Γ. This entanglement poses a problem because we saw in Section 3 that simulating the future impacts of climate change requires separately simulating both how climate change affects forecasts and how climate change affects the distribution of weather. In order to use an estimate of Γ to project the costs of future climate change, one has to believe that changes in the climate convey the same information as do realized changes in the weather: the $\partial\theta/\partial w$ terms embedded in Γ must adequately approximate $\partial\theta/\partial C$. However, the marginal effect of climate on a forecast θ may not be even roughly approximated by the marginal effect of weather.[28] If a known change in climate carries a stronger signal than a change in the weather, then Γ may underestimate long-run adaptation.

Burke and Emerick (2016) and Hsiang (2016) emphasize the value of comparing the $\hat{\beta}$ estimated via long differences to the β estimated from conventional time series regressions.

[27] In particular, if $\bar{w}$ is the average weather for the unit under consideration, then we have $Y_{t\to T}(\mathbf{w}) \to Y_{t\to T}(\bar{w})$ as $T \to \infty$ because the individual weather shocks average out to 0 in a stationary climate.

[28] As Dell et al. (2014, 779) observe when discussing long differences, adaptation depends on whether agents "perceived [the change in average temperature] to be a permanent change or just an accumulation of idiosyncratic shocks."

The authors argue that if these estimates are similar, then weather is a good proxy for climate, with long-run adaptation not playing an important role in the data. The present derivation clarifies what we learn when the estimated $\hat{\beta}$ and β are similar. Consider shrinking T towards t in equation (9). As we shrink T, we increase the relative contribution from the final line and eventually lose the second line altogether. The expectations terms change from being estimated in Γ to being a pure source of bias (via the time $t-1$ term, as described in Section 4). When the expectations channels are sufficiently small, then our estimated $\hat{\beta}$ will be comparable to our estimated β. In particular, these two estimates should be similar when the outcome Y_t does not depend on previous controls a_{t-1}. In this case, long-run adaptation is not important, as argued by Burke and Emerick (2016) and Hsiang (2016). However, these two estimates should also be similar when previous periods' durable investments are independent of previous weather, as when past weather does not contain a strong signal of future weather. In this case, we estimate similar β and $\hat{\beta}$ because weather shocks are not strongly informative about later weather. Critically, we have not learned that the climate is irrelevant for forecasts or adaptation: changing the climate is an experiment that plausibly carries different information than does the experiment of changing a period's weather.[29]

6 Conclusions

We have formally analyzed the implications of climate change for several types of outcomes in a dynamic setting that distinguishes the informational content of climate and weather. We have seen that climate change affects economic outcomes through direct weather channels and also through expectations channels. The recent empirical literature has focused on the direct weather channels, and I have described how future work may estimate the expectations channels. Further, we have also seen how ignoring expectations can bias estimates of the direct weather channels, which suggests a need to reevaluate the conclusions of recent empirical work. The net bias in projections of climate change impacts resulting from ignoring expectations channels and using biased weather channels is ambiguous in sign. Future empir-

[29] Burke and Emerick (2016) do find similar coefficients from the time series and long difference regressions, which they interpret as evidence of a lack of adaptation. They conduct an interesting set of checks to verify that their result is not due to agents' failure to recognize that the climate was changing: they show that farmers' responses do not depend on past experience of extreme weather, on the baseline variance of the weather, on education, or on political affiliation. These results suggest that farmers with less reason to extrapolate weather to climate performed the same type of extrapolation as did farmers with more reason to extrapolate weather to climate. However, these results are consistent not only with general recognition of climate change but also with a failure by all groups to extrapolate climate change from experienced weather variation. For the studied 1980–2000 period, it is perfectly plausible that $\partial\theta_t/\partial w_t \neq \partial\theta_t/\partial C$. See Kelly et al. (2005), Deryugina (2013), and Kala (2016) for more on learning about climate change from observations of the weather, see Bakkensen (2016) for a comparison of learning from personal observations and from official forecasts in the case of tornadoes, and see Libecap and Hansen (2002) for an analysis of learning about agricultural productivity from weather observations in the early twentieth century U.S.

ical work should seek variation in weather forecasts that can identify expectations channels and should seek truly exogenous variation in weather that can eliminate the potential bias in the standard approach to estimating the direct weather channels. Future work should also consider estimating structural models that can explicitly account for expectations and learning and allow for coherent simulation of climate counterfactuals.

Appendix: Proofs

Proof of Lemma 1

We have:

$$
\begin{aligned}
&\int \frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C}\,\mathrm{d}w_1\\
=&\int \left[\frac{\partial p_2(w_2;\theta_1)}{\partial \theta_1}\frac{\partial \theta_1(w_1, C)}{\partial C}\,p_1(w_1;C) + p_2(w_2;\theta_1)\,\frac{\partial p_1(w_1;C)}{\partial C}\right]\,\mathrm{d}w_1\\
=&\frac{\partial p_2(w_2;\theta_1)}{\partial \theta_1}\frac{\partial \theta_1(w_1, C)}{\partial C}\int p_1(w_1;C)\;\mathrm{d}w_1 + p_2(w_2;\theta_1)\int \frac{\partial p_1(w_1;C)}{\partial C}\;\mathrm{d}w_1\\
=&\frac{\partial p_2(w_2;\theta_1)}{\partial \theta_1}\frac{\partial \theta_1(w_1, C)}{\partial C},
\end{aligned}
$$

where we use Assumption 1 in the second equality and then recognize that probabilities integrate to 1, both before and after a marginal change in C. This proves the first part of the lemma.

We also have:

$$
\begin{aligned}
&\int\int \frac{\mathrm{d}p_0(w_1, w_2, w_3)}{\mathrm{d}C}\,\mathrm{d}w_2\,\mathrm{d}w_1\\
=&\frac{\partial p_3(w_3;\theta_2)}{\partial \theta_2}\frac{\partial \theta_2(w_2, C)}{\partial C}\int\int p_2(w_2;\theta_1)\,p_1(w_1;C)\,\mathrm{d}w_2\,\mathrm{d}w_1 + \int\int p_3(w_3;\theta_2)\,\frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C}\;\mathrm{d}w_2\,\mathrm{d}w_1\\
=&\frac{\partial p_3(w_3;\theta_2)}{\partial \theta_2}\frac{\partial \theta_2(w_2, C)}{\partial C} + \int p(w_3;\theta_2)\frac{\partial p_2(w_2;\theta_1)}{\partial \theta_1}\frac{\partial \theta_2(w_2, C)}{\partial C}\,\mathrm{d}w_2\\
=&\frac{\partial p_3(w_3;\theta_2)}{\partial \theta_2}\frac{\partial \theta_2(w_2, C)}{\partial C} + p_3(w_3;\theta_2)\int \frac{\mathrm{d}p_2(w_2;\theta_1(w_1, C))}{\mathrm{d}C}\,\mathrm{d}w_2\\
=&\frac{\partial p_3(w_3;\theta_2)}{\partial \theta_2}\frac{\partial \theta_2(w_2, C)}{\partial C},
\end{aligned}
$$

where we use Assumption 1 in the first equality, use the result from the first part of the lemma in the third line, and recognize in the last line that probabilities always integrate to 1. This proves the second part of the lemma.

Proof of Proposition 1 and Corollary 1

Applying the implicit function theorem to the first-order condition that defines $a_2^*(a_1, w_2, \theta_2)$, we have:

$$\frac{\partial a_2^*(a_1^*, w_2, \theta_2)}{\partial w_2} = -\frac{\frac{\partial^2 \pi_2}{\partial a_2 \partial w_2} + \beta \int \frac{\partial \pi_3}{\partial a_2} \frac{\partial p_3(w_3;\theta_2)}{\partial \theta_2} \frac{\partial \theta_2(w_2,C)}{\partial w_2} \, \mathrm{d}w_3}{\frac{\partial^2 \pi_2}{\partial a_2^2} + \beta E_2 \left[\frac{\partial^2 \pi_3}{\partial a_2^2} + \frac{\partial^2 \pi_3}{\partial a_2 \partial a_3} \frac{\partial a_3^*(a_2^*, w_3)}{\partial a_2} \right]}.$$

Using Assumption 2, we have:

$$\frac{\partial a_2^*(a_1^*, w_2, \theta_2)}{\partial w_2} = -\frac{\beta \frac{\partial \pi_3}{\partial a_2} \int \frac{\mathrm{d}p_3(w_3;\theta_2(w_2,C))}{\mathrm{d}w_2} \, \mathrm{d}w_3}{\frac{\partial^2 \pi_2}{\partial a_2^2} + \beta E_2 \left[\frac{\partial^2 \pi_3}{\partial a_2^2} + \frac{\partial^2 \pi_3}{\partial a_2 \partial a_3} \frac{\partial a_3^*(a_2^*, w_3)}{\partial a_2} \right]} = 0.$$

Similar analysis implies that $a_1^*(w_1, \theta_1)$ is independent of w_1.

Using the implicit function theorem on the first-order condition that defines $a_1^*(w_1, \theta_1)$, we have:

$$\frac{\partial a_1^*(w_1, \theta_1)}{\partial \theta_1} = -\frac{\beta \int \frac{\partial \pi_2}{\partial a_1} \frac{\partial p_2(w_2;\theta_1)}{\partial \theta_1} \, \mathrm{d}w_2}{\frac{\partial^2 \pi_1}{\partial a_1^2} + \beta E_1 \left[\frac{\partial^2 \pi_2}{\partial a_1^2} + \frac{\partial^2 \pi_2}{\partial a_1 \partial a_2} \frac{\partial a_2^*(a_1^*, w_2, \theta_2)}{\partial a_1} \right]}. \tag{A-1}$$

Assumption 2 implies that

$$\int \frac{\partial \pi_2}{\partial a_1} \frac{\partial p_2(w_2; \theta_1)}{\partial \theta_1} \, \mathrm{d}w_2 = \frac{\partial \pi_2}{\partial a_1} \int \frac{\partial p_2(w_2; \theta_1)}{\partial \theta_1} \, \mathrm{d}w_2 = 0,$$

and thus that $\partial a_1^* / \partial \theta_1 = 0$.

Using the implicit function theorem on the first-order condition that defines $a_2^*(a_1, w_2, \theta_2)$, we have:

$$\frac{\partial a_2^*(a_1^*, w_2, \theta_2)}{\partial \theta_2} = -\frac{\beta \int \frac{\partial \pi_3}{\partial a_2} \frac{\partial p_3(w_3;\theta_2)}{\partial \theta_2} \, \mathrm{d}w_3}{\frac{\partial^2 \pi_2}{\partial a_2^2} + \beta E_2 \left[\frac{\partial^2 \pi_3}{\partial a_2^2} + \frac{\partial^2 \pi_3}{\partial a_2 \partial a_3} \frac{\partial a_3^*(a_2^*, w_3)}{\partial a_2} \right]}. \tag{A-2}$$

Assumption 2 implies that

$$\int \frac{\partial \pi_3}{\partial a_2} \frac{\partial p_3(w_3; \theta_2)}{\partial \theta_2} \, \mathrm{d}w_3 = \frac{\partial \pi_3}{\partial a_2} \int \frac{\partial p_3(w_3; \theta_2)}{\partial \theta_2} \, \mathrm{d}w_3 = 0,$$

and thus that $\partial a_2^* / \partial \theta_2 = 0$.

Using the independence of a_t^* from w_t and θ_t, we have:

$$\int \int a_2^*(a_1^*, w_2, \theta_2) \frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C} \, \mathrm{d}w_2 \, \mathrm{d}w_1 = a_2^*(a_1^*, w_2, \theta_2) \int \int \frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C} \, \mathrm{d}w_2 \, \mathrm{d}w_1 = 0.$$

Using equation (3) and $\partial a_1^*/\partial\theta_1 = \partial a_2^*/\partial\theta_2 = 0$ establishes the first parts of the proposition and corollary.

Because w_{t-1} can only enter π_t via a_{t-1}^*, Assumption 2 implies that π_t is independent of w_{t-1}. We thus have:

$$\int\int \pi_2(a_1^*, a_2^*, w_2)\,\frac{\mathrm{d}p_0(w_1,w_2)}{\mathrm{d}C}\,\mathrm{d}w_2\,\mathrm{d}w_1 = \int \pi_2(a_1^*, a_2^*, w_2)\left(\int \frac{\mathrm{d}p_0(w_1,w_2)}{\mathrm{d}C}\,\mathrm{d}w_1\right)\mathrm{d}w_2.$$

The second part of the proposition follows from equation (2) and $\partial a_1^*/\partial\theta_1 = \partial a_2^*/\partial\theta_2 = 0$. Using Assumption 1 and Lemma 1 in the last expression yields:

$$\int\int \pi_2(a_1^*, a_2^*, w_2)\,\frac{\mathrm{d}p_0(w_1,w_2)}{\mathrm{d}C}\,\mathrm{d}w_2\,\mathrm{d}w_1 = \int \pi_2(a_1^*, a_2^*, w_2)\,\frac{\partial p_2(w_2;\theta_1)}{\partial\theta_1}\frac{\partial\theta_1(w_1,C)}{\partial C}\,\mathrm{d}w_2.$$

The second part of the corollary follows.

The third part of the proposition follows from equation (1), the foregoing analysis, and

$$\int\int\int \pi_3(a_2^*, a_3^*, w_3)\,\frac{\mathrm{d}p_0(w_1,w_2,w_3)}{\mathrm{d}C}\,\mathrm{d}w_3\,\mathrm{d}w_2\,\mathrm{d}w_1 = \int \pi_3(a_2^*, a_3^*, w_3)\left(\int\int \frac{\mathrm{d}p_0(w_1,w_2,w_3)}{\mathrm{d}C}\,\mathrm{d}w_1\,\mathrm{d}w_2\right)\mathrm{d}w_3$$

Using Assumption 1 and Lemma 1 in the last expression yields:

$$\int\int\int \pi_3(a_2^*, a_3^*, w_3)\,\frac{\mathrm{d}p_0(w_1,w_2,w_3)}{\mathrm{d}C}\,\mathrm{d}w_3\,\mathrm{d}w_2\,\mathrm{d}w_1 = \int \pi_3(a_2^*, a_3^*, w_3)\,\frac{\partial p_3(w_3;\theta_2)}{\partial\theta_2}\frac{\partial\theta_2(w_2,C)}{\partial C}\,\mathrm{d}w_3.$$

The third part of the corollary follows.

Proof of Proposition 2 and Corollary 2

From equations (A-1) and (A-2), Assumption 3 implies $\partial a_1^*/\partial\theta_1 = \partial a_2^*/\partial\theta_2 = 0$. Recognizing that π_2 being independent of a_1 implies that π_2 is independent of w_1, we have from equation (2):

$$\frac{\mathrm{d}E_0\left[\pi_2\right]}{\mathrm{d}C} = \int\int \pi_2(a_1^*, a_2^*, w_2)\,\frac{\mathrm{d}p_0(w_1,w_2)}{\mathrm{d}C}\,\mathrm{d}w_2\,\mathrm{d}w_1 = \int \pi_2(a_1^*, a_2^*, w_2)\left(\int \frac{\mathrm{d}p_0(w_1,w_2)}{\mathrm{d}C}\,\mathrm{d}w_1\right)\mathrm{d}w_2.$$

This establishes the first part of the proposition. Applying Assumption 1 and Lemma 1 then yields the first part of the corollary.

Now recognize that π_3 is independent of w_1 and w_2. We have:

$$\int\int\int \pi_3(a_2^*, a_3^*, w_3)\,\frac{\mathrm{d}p_0(w_1,w_2,w_3)}{\mathrm{d}C}\,\mathrm{d}w_3\,\mathrm{d}w_2\,\mathrm{d}w_1 = \int \pi_3(a_2^*, a_3^*, w_3)\left(\int\int \frac{\mathrm{d}p_0(w_1,w_2,w_3)}{\mathrm{d}C}\,\mathrm{d}w_1\,\mathrm{d}w_2\right)\mathrm{d}w_3$$

The second part of the proposition follows from the foregoing analysis and equation (1). Applying Assumption 1 and Lemma 1 then yields the second part of the corollary.

Applying the implicit function theorem to the first-order condition that defines a_2^*, we have:

$$\frac{\partial a_2^*(a_1^*, w_2, \theta_2)}{\partial a_1} = -\frac{\frac{\partial^2 \pi_2}{\partial a_2 \partial a_1}}{\frac{\partial^2 \pi_2}{\partial a_2^2} + \beta E_2 \left[\frac{\partial^2 \pi_3}{\partial a_2^2} + \frac{\partial^2 \pi_3}{\partial a_2 \partial a_3} \frac{\partial a_3^*(a_2^*, w_3)}{\partial a_2} \right]}. \tag{A-3}$$

This expression is zero under Assumption 3. Because a_2^* is independent of a_1, it is also independent of w_1. Recognizing once again that $\partial a_1^*/\partial \theta_1 = \partial a_2^*/\partial \theta_2 = 0$, we have from equation (3):

$$\begin{aligned} \frac{\mathrm{d}E_0\left[a_2^*\right]}{\mathrm{d}C} &= \int \int a_2^*(a_1^*, w_2, \theta_2) \frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C} \, \mathrm{d}w_2 \, \mathrm{d}w_1 \\ &= \int a_2^*(a_1^*, w_2, \theta_2) \left(\int \frac{\mathrm{d}p_0(w_1, w_2)}{\mathrm{d}C} \, \mathrm{d}w_1 \right) \mathrm{d}w_2. \end{aligned}$$

This establishes the third part of the proposition. Applying Assumption 1 and Lemma 1 then yields the third part of the corollary.

References

Ackerman, Frank, Stephen DeCanio, Richard Howarth, and Kristen Sheeran (2009) “Limitations of integrated assessment models of climate change,” *Climatic Change*, Vol. 95, No. 3-4, pp. 297–315.

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.

Auffhammer, Maximilian, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel (2013) “Using weather data and climate model output in economic analyses of climate change,” *Review of Environmental Economics and Policy*, Vol. 7, No. 2, pp. 181–198.

Auffhammer, Maximilian and Erin T. Mansur (2014) “Measuring climatic impacts on energy consumption: A review of the empirical literature,” *Energy Economics*, Vol. 46, pp. 522–530.

Bakkensen, Laura (2016) “Public versus private risk information: Evidence from U.S. tornadoes,” working paper, University of Arizona.

Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro (2016) “Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century,” *Journal of Political Economy*, Vol. 124, No. 1, pp. 105–159.

Barreca, Alan I. (2012) "Climate change, humidity, and mortality in the United States," *Journal of Environmental Economics and Management*, Vol. 63, No. 1, pp. 19–34.

Baylis, Patrick (2017) "Temperature and temperament," working paper.

Burke, Marshall, John Dykema, David B. Lobell, Edward Miguel, and Shanker Satyanath (2014) "Incorporating climate uncertainty into estimates of climate change impacts," *Review of Economics and Statistics*, Vol. 97, No. 2, pp. 461–471.

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

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

Cachon, Gerard, Santiago Gallino, and Marcelo Olivares (2012) "Severe weather and automobile assembly productivity," Research Paper 12/37, Columbia Business School.

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

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 (2013) "How do people update? The effects of local weather fluctuations on beliefs about global warming," *Climatic Change*, Vol. 118, No. 2, pp. 397–416.

Deryugina, Tatyana and Solomon M. Hsiang (2014) "Does the environment still matter? Daily temperature and income in the United States," Working Paper 20750, 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.

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.

Hill, Harvey S. J. and James W. Mjelde (2002) "Challenges and opportunities provided by seasonal climate forecasts: A literature review," *Journal of Agricultural and Applied Economics*, Vol. 34, No. 3, pp. 603–632.

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

Jia, Liwei, Xiaosong Yang, Gabriel A. Vecchi, Richard G. Gudgel, Thomas L. Delworth, Anthony Rosati, William F. Stern, Andrew T. Wittenberg, Lakshmi Krishnamurthy, Shaoqing Zhang, Rym Msadek, Sarah Kapnick, Seth Underwood, Fanrong Zeng, Whit G. Anderson, Venkatramani Balaji, and Keith Dixon (2014) "Improved seasonal prediction of temperature and precipitation over land in a high-resolution GFDL climate model," *Journal of Climate*, Vol. 28, No. 5, pp. 2044–2062.

Kala, Namrata (2016) "Learning, adaptation, and climate uncertainty: Evidence from Indian agriculture," working paper.

Katz, Richard W. and Jeffrey K. Lazo (2011) "Economic value of weather and climate forecasts," in Michael P. Clements and David F. Hendry eds. *The Oxford Handbook of Economic Forecasting*, Oxford: Oxford University Press.

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.

Kirtman, B., S. B. Power, J. A. Adedoyin, G. J. Boer, R. Bojariu, I. Camilloni, F. J. Doblas-Reyes, A. M. Fiore, M. Kimoto, G. A. Meehl, M. Prather, A. Sarr, C. Schar, R. Sutton, G. J. van Oldenborgh, G. Vecchi, and H. J. Wang (2013) "Near-term climate change: Projections and predictability," in T. F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. M. Midgley eds. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, UK and New York, NY, USA: Cambridge University Press, pp. 953–1028.

Lave, Lester B. (1963) "The value of better weather information to the raisin industry," *Econometrica*, Vol. 31, No. 1/2, pp. 151–164.

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.

Libecap, Gary D. and Zeynep Kocabiyik Hansen (2002) ""Rain follows the plow" and dry-farming doctrine: The climate information problem and homestead failure in the upper Great Plains, 1890–1925," *The Journal of Economic History*, Vol. 62, No. 1, pp. 86–120.

Melillo, Jerry M., Terese (T.C.) Richmond, and Gary W. Yohe eds. (2014) *Climate Change Impacts in the United States: The Third National Climate Assessment*: U.S. Global Change Research Program.

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.

Meza, Francisco J., James W. Hansen, and Daniel Osgood (2008) "Economic value of seasonal climate forecasts for agriculture: review of ex-ante assessments and recommendations for future research," *Journal of Applied Meteorology and Climatology*, Vol. 47, No. 5, pp. 1269–1286.

Miller, Benjamin (2015) "Does validity fall from the sky? Observant farmers and the endogeneity of rainfall," in *Dissertation: Weather, Expectations, and Complex Incentives*: University of California, San Diego, pp. 1–79.

Mjelde, James W., John B. Penson, and Clair J. Nixon (2000) "Dynamic aspects of the impact of the use of perfect climate forecasts in the corn belt region," *Journal of Applied Meteorology*, Vol. 39, No. 1, pp. 67–79.

Neidell, Matthew (2009) "Information, avoidance behavior, and health: The effect of ozone on asthma hospitalizations," *Journal of Human Resources*, Vol. 44, No. 2, pp. 450–478.

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

Pindyck, Robert S. (2013) "The climate policy dilemma," *Review of Environmental Economics and Policy*, Vol. 7, No. 2, pp. 219–237.

Ranson, Matthew (2014) "Crime, weather, and climate change," *Journal of Environmental Economics and Management*, Vol. 67, No. 3, pp. 274–302.

Roll, Richard (1984) "Orange juice and weather," *The American Economic Review*, Vol. 74, No. 5, pp. 861–880.

Rosenzweig, Mark R. and Christopher Udry (2014) "Rainfall forecasts, weather, and wages over the agricultural production cycle," *American Economic Review: Papers and Proceedings*, Vol. 104, No. 5, pp. 278–283.

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

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

Schlenker, Wolfram and Michael J. Roberts (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 Deschenes (2016) "A forward looking Ricardian approach: Do land markets capitalize climate change forecasts?" Working Paper 22413, National Bureau of Economic Research.

Shrader, Jeffrey (2017) "Expectations and adaptation to environmental risks," job market paper, University of California San Diego.

Smith, Doug M., Stephen Cusack, Andrew W. Colman, Chris K. Folland, Glen R. Harris, and James M. Murphy (2007) "Improved surface temperature prediction for the coming decade from a global climate model," *Science*, Vol. 317, No. 5839, pp. 796–799.

Solow, Andrew R., Richard F. Adams, Kelly J. Bryant, David M. Legler, James J. O'Brien, Bruce A. McCarl, William Nayda, and Rodney Weiher (1998) "The value of improved ENSO prediction to U.S. agriculture," *Climatic Change*, Vol. 39, No. 1, pp. 47–60.

Wood, Stephen A., Amir S. Jina, Meha Jain, Patti Kristjanson, and Ruth S. DeFries (2014) "Smallholder farmer cropping decisions related to climate variability across multiple regions," *Global Environmental Change*, Vol. 25, pp. 163–172.

Yang, Xiaosong, Gabriel A. Vecchi, Rich G. Gudgel, Thomas L. Delworth, Shaoqing Zhang, Anthony Rosati, Liwei Jia, William F. Stern, Andrew T. Wittenberg, Sarah Kapnick, Rym Msadek, Seth D. Underwood, Fanrong Zeng, Whit Anderson, and Venkatramani Balaji (2015) "Seasonal predictability of extratropical storm tracks in GFDL's high-resolution climate prediction model," *Journal of Climate*, Vol. 28, No. 9, pp. 3592–3611.

Zhang, Peng, Junjie Zhang, Olivier Deschenes, and Kyle C. Meng (2016) "Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants," working paper.