LEARNING, ADAPTATION, AND WEATHER IN A CHANGING CLIMATE*

FRANCES C. MOORE

Environmental Science and Policy

University of California Davis, 2140 Wickson Hall

One Shields Avenue, Davis, CA 95616, USA

fmoore@ucdavis.edu

Published 21 August 2017

Climate change will push the weather experienced by people affected outside the bounds of historic norms, resulting in unprecedented weather events. But people and firms should be able to learn from their experience of unusual weather and adjust their expectations about the climate distribution accordingly. The efficiency of this learning process gives an upper bound on the rate at which adaptation can occur and is therefore important in determining the adjustment costs associated with climate change. Learning about climate change requires people to infer the state of a changing probability distribution (climate) given annual draws from that distribution (weather). If the climate is stationary, it can be inferred from the distribution of historic weather observations, but if it is changing, the inference problem is more challenging. This paper first develops different learning models, including an efficient hierarchical Bayesian model in which the observer learns whether the climate is changing and, if it is, the functional form that describes that change. I contrast this with a less efficient but simpler learning model in which observers react to past changes but are unable to anticipate future changes. I propose a general metric of learning costs based on the average, discounted squared difference between beliefs and the true climate state and use climate model output to calculate this metric for two emissions scenarios, finding substantial relative differences between learning models and scenarios but small absolute values. Geographic differences arise from spatial patterns of warming rates and natural weather variability (noise). Finally, I present results from an experimental game simulating the adaptation decision, which suggests that people are able to learn about a trending climate and respond proactively.

Keywords: Adaptation; learning; weather; climate; extremes; Bayesian learning.

1. Introduction

Much climate change adaptation is a private good and will therefore be undertaken by individuals and firms acting in their own self-interest (Mendelsohn, 2000). However, adapting to climate change is not straightforward because it requires people to learn

*This paper contains supplementary material available on the journal website. The supplementary information includes details on Learning, Adaptation, and Weather in a Changing Climate.

This is an Open Access article published by World Scientific Publishing Company. It is distributed under the terms of the Creative Commons Attribution 4.0 (CC-BY) License. Further distribution of this work is permitted, provided the original work is properly cited.

about a changing climate and, in some cases, to anticipate future changes. The rate at which people perceive and learn about a changing climate places an upper limit on the rate of adaptation.

Learning about climate change is challenging because climate — the probability distribution from which weather realizations are drawn — is not directly observed. Instead, it must be inferred by agents based on the information available to them. In a stationary climate, the parameters of the climate distribution can be reliably inferred from historic weather observations. But if the climate is changing, then past observations are not necessarily informative about the current state of the probability distribution over possible weather outcomes. It seems unlikely that climate models will be able to provide the local-scale, specific, reliable, and trusted information required for most private adaptation decisions. Instead, direct experience of weather events is likely to be an important information source, with several studies suggesting a link between experience of weather anomalies and stated belief in climate change (Howe *et al.*, 2014; McCright *et al.*, 2014; Kaufmann *et al.*, 2017). Therefore, how well people learn from experience in nonstationary, stochastic environments has important implications for the rate of adaptation and the economic costs associated with climate change.

The rate at which private adaptation occurs is in turn an important determinant of adjustment costs and of the overall economic costs associated with climate change. A change in climate has both equilibrium effects (differences in productivity after full adjustment to the new climate) and adjustment costs (costs associated with being imperfectly adjusted to the new climate). Adjustment costs arise because it takes time to adopt the full set of possible adaptation options. This lag may be due to learning rates, long turnover time of capital stock, or market failures such as a lack of information about available adaptation options (Kelly *et al.*, 2005). Several authors have argued that, because climate change happens gradually, costs associated with the turnover of capital stock are likely to be small (Mendelsohn *et al.*, 1994; Kelly *et al.*, 2005). The rate of learning may therefore be a principal determinant of climate change adjustment costs, which in turn have important implications for both mitigation and adaptation policy (Cropper and Oates, 1992; Stern, 2006). The goal of this paper is to begin to provide an estimate of the magnitude of global adjustment costs associated with learning by combining learning models with climate model predictions and some experimental evidence.

This study builds on previous work by Kelly, Kolstad, and Mitchell (KKM) (Kelly *et al.*, 2005), who first proposed using a Bayesian updating model to describe how agents learn about climate change, given observations of weather in order to quantify dynamic adjustment costs. They use a Bayesian updating model in which the agent updates beliefs about a set of parameters describing the time evolution of the mean state of the climate using sequential weather observations, starting with priors informed by the IPCC. Using an empirical example from U.S. agriculture,

KKM find that learning this way takes decades and imposes significant adjustment costs.

This paper develops the ideas in KKM in two ways. First, it extends the Bayesian learning model using a more flexible hierarchical model. This allows for the fact that agents are unlikely to know *a priori* the particular functional form describing how climate changes over time. Instead, they must use weather observations to jointly infer *whether* the climate is changing, *when* it started changing, and the *rate* at which it is changing. The particular hierarchical Bayesian model described here covers a broad class of functional forms relevant to climate change in which the climate is initially stationary and begins changing at a certain point in time. The agent integrates over all possible functional forms in order to determine the current probability distribution over weather outcomes. The KKM model is a particular case of this model in which the agent places zero probability on the possibility that the climate is stationary and has priors over the rate of change informed by the IPCC.

Secondly, this paper investigates the importance of *how* observers learn about climate change for determining adjustment costs. Though Bayesian learning is efficient, it can be computationally burdensome. In reality, people may be updating beliefs using less efficient but more straightforward heuristics (Simon, 1955; Kahneman, 2003). In addition, both KKM and the hierarchical Bayesian model presented in this paper model a process of proactive learning and adaptation — agents are able to learn that the climate is changing and therefore anticipate and prepare for future change. In reality, much adaptation may instead be reactive — agents are able to learn and adjust to past changes, but not anticipate future changes. Many formalized updating rules in institutions are reactive (for instance, zoning rules based on the historic 50- or 100-year flood) and therefore will always be suboptimal if the climate is not stationary. This paper compares the efficient, proactive hierarchical Bayesian learning model with a simpler but reactive rolling-mean learning model in order to assess the importance of learning efficiency in determining adjustment costs. Finally, I present the results of a simple experiment designed to distinguish between proactive and reactive learning models in a laboratory context. This provides evidence that agents are able to learn about a changing climate and to anticipate future changes and respond proactively in a manner consistent with the Bayesian learning model.

This paper contributes to the existing literature in two ways. Topically, this study contributes to a small but growing body of work on the rate of adaptation to climate change. To date, these studies have empirically estimated adjustment rates in response to historical climate or weather changes in particular locations and sectors. Hornbeck (2012) looked at responses to productivity shocks associated with the Dust Bowl and found that it took decades for farms to fully adjust. Taraz, (in press) uses decadal changes in the Indian monsoon to show that farmers are able to adapt to changes on the 10–20-year timescale, although this adaptation is imperfect and significant residual impacts from climate fluctuations remain. Burke and Emerick (2016) use the response

of U.S. maize yields to multi-decadal trends in growing-season temperature to show limited ability of farmers to adapt to changing climate conditions in the medium run. Kala (2016) is perhaps most similar to this paper because it focusses on how farmers use observations of the monsoon start date to update their beliefs about the underlying climate distribution. She uses observations of planting decisions to test alternative learning models, finding evidence for ambiguity aversion in the way farmers adjust their beliefs in response to weather experiences. Although not explicitly discussing adjustment costs, Kelly and Kolstad (1999) model Bayesian learning about the climate (specifically about the climate sensitivity) using weather observations, finding that learning is challenging and takes several decades. This paper extends this body of work by describing a general theoretical framework for learning-related adjustment costs and connecting a measure of these costs to global rates of climate change using climate model output. It is therefore linking the rate of learning in individual regions and sectors studied by previous papers, the global adjustment costs associated with climate change, and the rate of climate change as determined by global mitigation decisions.

Theoretically, this paper contributes to a literature on the economic consequences of heuristic, imperfectly rational decision-making by economic agents. There is now a large literature showing that agents do not always behave rationally and, in particular, that they show significant biases in processing of probabilistic information (Tversky and Kahneman, 1974; Gilovich *et al.*, 1985; Kahneman *et al.*, 1991; Kahneman, 2003). However, the aggregate economic implications of this suboptimal decision-making are less well understood. Akerlof and Yellen (1985) show that though the envelope theorem means that in many cases the economic consequences of small deviations from rationality may be small for individuals, the aggregate effects on the economy could be an order of magnitude larger. Klumpp (2006) compares decisions under an efficient Bayesian learning model and a heuristic linear learning model in a simulated environment. He finds that although different learning models result in different beliefs, if the decision space is discrete, these beliefs could result in identical or near-identical decisions. More generally, the welfare consequences of imperfect learning models will depend on both the difference in beliefs resulting from alternative learning models and the nature of the decision space (in this case adaptation options) available to agents. This paper addresses the former — quantifying the expected difference in beliefs under different learning models and climate scenarios.

The following section describes two models for how agents might infer the state of the climate-given observations of weather. Section 3 describes how beliefs are connected to adjustment costs and simulates beliefs and associated learning costs over the globe for two different emission scenarios under alternative learning models. Section 4 presents results from a simple experiment designed to distinguish between learning models in a laboratory setting. Section 5 concludes.

2. Learning Models

Weather at time t is a draw from the climate distribution at time t . Here the climate distribution is assumed to be normal, with a mean that may be changing over time¹:

$$w_t \sim N(\mu_t, \sigma_w^2).$$

The initial baseline climate is known:

$$w_0 \sim N(\mu_0, \sigma_w^2).$$

At some point, r , the mean begins changing according to a functional form described by a set of parameters, β . Therefore, at time t , the mean state of the climate is given by:

$$\mu_t = \begin{cases} \mu_0 & \text{for } t \leq r, \\ f(\beta, t) & \text{for } t > r. \end{cases} \quad (1)$$

The natural variability of weather around the mean (σ_w^2) does not change, and I assume that the learner knows both the baseline climate and that the natural weather variability is not changing. The assumption of unchanging weather variability is consistent with some recent evidence for global surface temperature, which suggests that once the nonstationarity of the climate mean is accounted for, there is not strong evidence in observations or models for increasing variability around the mean (Rhines and Huybers, 2013; Thompson *et al.*, 2015). This would not hold for precipitation, the variability of which is projected to increase with climate change (Marvel and Bonfils, 2013). Relaxing the assumption of known and unchanging weather variability would substantially increase the complexity of the learning models and is not attempted in this paper. The function f is arbitrary and, for simplicity I assume a linear rate of change so that $f(\beta, t) = \mu_0 + \beta(t - r)$ for $t > r$, as in KKM.

The observer's problem at time t is to infer, given a time series of weather observations up to that point, w_t , the current distribution over possible weather outcomes that they face. In other words, they are learning the predictive distribution of future weather outcomes, conditional on the weather that has been observed to date $P(w_{t+1}|w_t)$.² The learner sequentially updates their beliefs about the state of the climate, given each new weather observation and a particular learning model. The time series describing the learner's beliefs about the state (mean and variance) of the climate system therefore depends on how the climate is changing, the particular weather realizations that the learner observes from the changing climate, and the learning model they are using to update beliefs. I examine two learning models in this paper.

¹A normal distribution is reasonable distribution for annual temperatures in most locations, but may not be a good fit for other climate variables such as rainfall, which tend to be more heavily skewed (Wilson and Toumi, 2005).

²The posterior predictive distribution, $P(w_{t+1}|w_t)$, in this paper, is defined only for the period immediately following t , not for subsequent periods. Therefore, this framework is most directly relevant for adaptation decisions made on a year-to-year basis (for instance, changing crop varieties or planting dates in agriculture), rather than capital investment decisions. For large capital investments that last multiple years, agents will use their beliefs over multiple future years to inform their decision (i.e., $P(w_{t+1}w_{t+2} \dots w_{t+n}|w_t)$, where n is the lifetime of the investment, something not directly addressed in this paper.

2.1. Proactive learning: Hierarchical Bayesian model

Bayesian updating is an efficient way of incorporating new information into existing beliefs. As long as the learner places a nonzero prior probability on the true model, beliefs are guaranteed to eventually converge to the true state. Although standard Bayesian models are well known, the hierarchical model described here allows for the fact that observers may be unsure whether or not the climate is changing and, if it is changing, the functional form that describes how it is changing. The Bayesian learner in this model must infer whether or not the climate is changing, when it started changing, and how quickly it is changing. In other words, he uses the observed weather ($\mathbf{w}_t$) to update beliefs over the probability of both the change point (r) and possible rates of change (β), integrating over these posterior distributions in order to give the posterior predictive distribution, $P(\mathbf{w}_{t+1}|\mathbf{w}_t)$ (Adams and MacKay, 2007).³

The probability of a weather event, $\mathbf{w}_{t+1}$, given the history of weather experienced to date, $\mathbf{w}_t$, is the probability of that event under each possible change-point, multiplied by the probability of that change-point occurring:

$$P(\mathbf{w}_{t+1}|\mathbf{w}_t) = \sum_{n=1}^t P(\mathbf{w}_{t+1}|r_n, \mathbf{w}_t) * P(r_n|\mathbf{w}_t). \quad (2)$$

The second term in Eq. (2) is the posterior probability density over possible change-points, which comes from Bayes law:

$$P(r_n|\mathbf{w}_t) = \frac{P(r_n) * P(\mathbf{w}_t|r_n)}{P(\mathbf{w}_t)}.$$

$P(r_n)$ is the prior probability of a change-point occurring at point n whereas $P(\mathbf{w}_t|r_n)$ is the likelihood of the data, conditional on a particular change-point. Obtaining this likelihood requires integrating over all possible rates of change:

$$P(\mathbf{w}_t|r_n) = \int_{-\infty}^{\infty} P(\mathbf{w}_t|r_n, \beta) * P(\beta|r_n) d\beta. \quad (3)$$

$P(\beta|r_n)$ is the prior probability over the rate of change, conditional on a particular change-point. I assume that this prior is normal, distributed $N(\mu_{\beta, r}, \sigma_{\beta, r}^2)$. Given a particular rate of change and change-point, the probability of the set of weather observations is also normal $N(\mu_t, \sigma_w^2)$, where μ_t is defined according to Eq. (1). Under these distributional assumptions, Eq. (3) is a Gaussian integral with an exact solution given in Appendix A.

The first term in Eq. (2), the probability of a particular weather observation given a change-point and set of weather observations, requires integrating over all possible

³The model could, with some difficulty, be extended so that observers chose between a set of alternative functional forms describing the rate of climate change (i.e., $f_1(\beta_1, t)$, $f_2(\beta_2, t)$ etc.).

rates of change:

$$P(w_{t+1}|r_n, \mathbf{w}_t) = \int_{-\infty}^{\infty} P(w_{t+1}|r_n, \beta) * P(\beta|r_n, \mathbf{w}_t) d\beta.$$

The second term in this integral is simply the posterior distribution over rates of change and, given that both the prior and the likelihood are normally distributed, is also normally distributed (SI Appendix). The first term is also normally distributed ($N(\mu_{t+1}\sigma_w^2)$, with μ_{t+1} defined according to Eq. (1)), meaning that this is also a Gaussian integral with an exact solution given in Appendix A. The posterior predictive distribution $P(w_{t+1}|\mathbf{w}_t)$, although not itself a standard distribution, can be sampled using the two Gaussian integrals (additional details in Appendix A).

As with any Bayesian model, priors must be specified and form the free parameters of the model. Priors are sequentially updated with each weather observation, so that the posterior for one period becomes the prior for the subsequent period:

$$P(r_n)_{t+1} = \frac{P(r_n)_t * P(\mathbf{w}_t|r_n)}{P(\mathbf{w}_t)},$$

$$P(\beta|r_n)_{t+1} = \frac{P(\beta|r_n)_t * P(\mathbf{w}_t|\beta, r_n)}{P(\mathbf{w}_t|r_n)},$$

where $P(r_n)_t$ is the prior belief that a change-point occurred at time n after t observations and $P(\beta|r_n)_t$ is the prior belief at period t over the rate of change, conditional on a change-point occurring at time n .

2.2. Reactive learning: Rolling-mean model

Although Bayesian learning efficiently uses new information and is guaranteed to converge to the true state, it is computationally very complex. It may be that people will instead use a much simpler process to learn about climate change. In addition to computational complexity, the Bayesian model assumes proactive learning — learners are able to perceive that the climate is nonstationary and therefore to anticipate future changes. In reality, it may be that people are able to react to climate change that has happened, but not anticipate future changes. In particular, many institutional updating rules relevant to climate change adaptation are reactive. For example, the yield basis for crop insurance in the United States is typically calculated using a 10-year trailing mean (Plastina and Edwards) or the definition of a floodplain or storm drain design specifications might be based on historical 100-year flood events. While these rules do result in gradual adaptation to past change, because they do not explicitly recognize the nonstationary nature of climate change, they will never produce beliefs that converge to the true climate and therefore will always result in suboptimal decisions under climate change.

Here I use the rolling-mean process as an example of a computationally simple, reactive learning process. The mean of the climate distribution at time t is estimated by

averaging a fixed number of past weather observations:

$$\widehat{\mu}_{t+1} = \frac{\sum_{t-(d-1)}^t w_t}{d},$$

where d is the window over which the learner considers observations in their estimate of the next mean state $\widehat{\mu}_{t+1}$. Given this belief about the mean climate state, the learner assumes that weather observations are distributed about the mean with natural weather variability σ_w^2 . So the predictive distribution for the rolling-mean learner is given by the normal distribution $P(w_{t+1}|\mathbf{w}_t) = N(\widehat{\mu}_{t+1}, \sigma_w^2)$.

Figure 1 shows the behavior of these two learning models under a changing climate, as well as a reference case without learning. If there is no learning, then the climate state quickly departs from the baseline, resulting in weather events that are consistently different from expectations (gray lines, Fig. 1). If instead the agent is able to use observations of weather to update their beliefs about the state of the climate, then beliefs are able to more closely track the changing climate state (pink and blue lines, Fig. 1).

In both learning models, beliefs lag the true state of the climate for some period after the climate begins changing because learners have yet to acquire sufficient new weather observations to conclude that a change has occurred. Over time, both types of learners make new observations of weather and conclude that the climate is warming. The Bayesian updater is able to learn about the rate at which the climate is changing

Figure 1. Performance of the two learning models presented in this paper and a reference case with no learning under a stylized representation of a changing climate. Solid lines show beliefs about the mean state of the climate and dashed lines show ± 2 standard deviations of the perceived weather distribution based on 250 simulations

and his beliefs therefore ultimately converge to the true state of the climate.⁴ The rolling-mean learner, in contrast, never learns about the rate of climate change. Because of this, he is not able to anticipate future change and his beliefs never converge to the true climate state.

3. Learning and Adjustment Costs

3.1. General framework

The fact that climate change is imperfectly observed by those affected results in a difference between beliefs about the climate and its true state. The observer anticipates and prepares for weather from the distribution $P(w_{t+1}|\mathbf{w}_t)$ when, in fact, weather events are drawn from the true climate distribution, $N(\mu_{t+1}, \sigma_w^2)$. The difference between these distributions means that the observer experiences a set of weather events that is different from what they expect.

Adjustment costs arise if the observer would have done something different in response to the actual climate distribution relative to the climate distribution they expect. Following KKM, if observers face a profit function, y , that depends on weather and a set of management parameters θ , they will choose θ to maximize profits, conditional on beliefs about the climate distribution:

$$\theta^*(P(w_{t+1}|\mathbf{w}_t)) = \max_{\theta} \int_{\underline{w}}^{\bar{w}} y(w_{t+1}, \theta) * P(w_{t+1}|\mathbf{w}_t) dw_{t+1},$$

where $\bar{w}$ and $\underline{w}$ are plausible upper and lower bounds of the climate distribution. Adjustment costs at time t (C_t) are the difference in expected profits given chosen management options ($\theta^*(P(w_{t+1}|\mathbf{w}_t))$) and the options that would have been chosen with full information about the climate state ($\theta^*(N(\mu_{t+1}, \sigma_w^2))$), given that weather is actually drawn from the true climate distribution $N(\mu_{t+1}, \sigma_w^2)$. They will differ depending on the agent's (a) production function and management options, on the updating rule (u) used to produce the posterior predictive distribution, and on the climate mean, variability, and weather draws in a particular location (l):

$$\begin{aligned} C_{tula} = & \int_{\underline{w}}^{\bar{w}} y_a(w_{l,t+1}, \theta_a^*(N(\mu_{l,t+1}, \sigma_{w,l}^2))) * N(\mu_{l,t+1}, \sigma_{w,l}^2) dw_{l,t+1} \\ & - \int_{\underline{w}}^{\bar{w}} y_a(w_{l,t+1}, \theta_a^*(P_u(w_{l,t+1}|\mathbf{w}_{1,t}))) * N(\mu_{l,t+1}, \sigma_{w,l}^2) dw_{l,t+1}. \end{aligned}$$

⁴The KKM model forms one part of the hierarchical Bayesian model — if agents knew for sure that climate was changing, when it began changing, and held priors informed by the IPCC, results would be identical. The principal contribution of the hierarchical Bayesian model is in modeling how agents learn the functional form describing how the climate is changing. Once this is learned, beliefs under the two models will converge so that, in the long run, beliefs in both the KKM model and the hierarchical Bayesian model will converge to the true climate state.

A necessary and sufficient condition for the existence of adjustment costs for a particular type of agent in a particular place is that $\theta_a^*(P_u(w_{l,t+1}|\mathbf{w}_{l,t})) \neq \theta_a^*(N(\mu_{l,t+1}, \sigma_{w,l}^2))$. This requires both that beliefs about the climate state are inaccurate (i.e., $P_u(w_{l,t+1}|\mathbf{w}_{l,t}) \neq N(\mu_{l,t+1}, \sigma_{w,l}^2)$) and that optimal management choices are different under the expected and real climate distribution.

The present value of the full set of learning-related adjustment costs associated with a particular updating rule under climate change requires summing adjustment costs from all agents and locations and discounting back to the present:

$$C_u = \sum_{t=0}^{\infty} \sum_a \sum_l \frac{C_{ula}}{(1+r)^{t+1}}.$$

Precisely determining the global value of C_u requires knowing the production function and adaptation options for each agent affected by climate change, something that is currently unrealistic. Therefore, in order to evaluate the performance of different updating rules, I abstract away from the production function and instead assume that adjustment costs will be weakly increasing in the difference between the perceived and true climate means, based on a loss function L . This gives a quasi-adjustment cost metric ($\tilde{C}$) that depends on the location and updating rule:

$$\tilde{C}_{ul} = \sum_{t=0}^{\infty} \frac{L(\mu_{l,t+1} - E(P_u(w_{l,t+1}|\mathbf{w}_{l,t})))}{(1+r)^{t+1}}.$$

These quasi-adjustment costs are aggregated over regions using a population-weighted averaging:

$$\tilde{C}_u = \frac{\sum_l N_l \sum_{t=0}^{\infty} \frac{L(\mu_{l,t+1} - E(P_u(w_{l,t+1}|\mathbf{w}_{l,t})))}{(1+r)^{t+1}}}{\sum_l N_l},$$

where N_l is the population in location l . The loss function L is arbitrary and for the simulations in this paper I use a discounted root mean squared error measure over a finite time horizon, so that $\tilde{C}_u$ becomes:

$$\tilde{C}_u = \frac{\sum_l N_l \sqrt{\frac{1}{T} \sum_{t=0}^T \frac{(\mu_{l,t+1} - E(P_u(w_{l,t+1}|\mathbf{w}_{l,t})))^2}{(1+r)^{t+1}}}}{\sum_l N_l}.$$

The benefits of this metric are that the units are the same as the relevant climate variable and therefore interpretable in absolute, not just in relative, terms and that the squared term accounts for the fact that larger deviations are likely to be relatively more damaging than smaller ones. The interpretation of this metric is as the population weighted-average difference between the perceived and the true climatological means, adjusting for time preferences, and for the fact that larger differences are likely to be more damaging than small differences.

3.2. Learning costs for two emissions scenarios

To show how the quasi-adjustment cost metric $\tilde{C}_u$ varies depending on the updating model and future emissions scenarios, I calculate $\tilde{C}_u$ over the period 2010–2050 under both a business-as-usual scenario (RCP 8.5) and an aggressive mitigation scenario (RCP 2.6), focussing on annual average temperatures as the weather metric of interest. Surface temperature trends come from the CMIP5 model ensemble of 81 runs of 39 different climate models, available at 0.5° resolution. For each land grid cell, I calculate the linear trend in temperature between 2010 and 2050 (Fig. A.1(a)). Natural inter-annual variability of temperature comes from a 500-year (800AD–1300AD) pre-industrial control run of a single climate model (the Community Climate System Model 4, Fig. A.1(b)). For computational reasons, I bin each grid cell into one of 16 groups based on the combination of climate trend and internal variability in that grid cell. For each group, I simulate 100 possible weather realizations based on an initialization of 10 years of a stationary climate followed by 40 years of linear change at the rate estimated for that group from the CMIP5 ensemble mean.

Both learning processes include some free parameters. For the Bayesian learner, priors over the rate of change ($P(\beta|r_n)$) are centered fairly tightly around zero ($N(0, 0.001)$). Priors over the location of the change-point ($P(r_n)$) are uniform and are updated sequentially with additional weather observations. The Bayesian also holds a fixed belief that the probability of a change-point occurring at all is 20%, which is not updated. The rolling-mean process has a single free parameter — the window over which past observations are averaged to give the current estimate of the climate state. This is set at 15 years in the simulations. The discount rate, r , is 2.5%.

Figure 2(a) gives the average, population-weighted learning-related adjustment costs for the two learning models and the no-learning reference case for the period 2010–2050 for the two emissions scenarios. As expected, the learning costs are greater than zero and decrease with the efficiency of the learning model. Costs in the no-learning case are 2.9–3.1 times higher than the rolling-mean learner and 4.5–5.5 times higher than the Bayesian learner. Figure 2(a) also suggests an interaction between the rate of climate change and the benefits of learning more efficiently. Moving from moderate rates of change under RCP 2.6 to faster climate change under RCP 8.5 increases learning costs uniformly, but this increase is less (84%) for the Bayesian learner than for the rolling-mean learner (116%).

Note that the discount rate is likely to be important in determining the relative importance of adjustment costs under different learning models, as well as the absolute value of $\tilde{C}_u$. As can be seen in Fig. 1, learning models differ most in later time periods. Initially, differences between beliefs and the true climate state are small (because the climate has not yet changed significantly). The difference grows over time as the climate continues changing, but there is not yet sufficient evidence, in the form of altered weather conditions, for learners to adjust their beliefs. Only after new information has been acquired, in the form of unexpected weather conditions, do learners

begin adjusting their beliefs and start to diverge from nonlearners. The distinction between the rolling-mean learner and the Bayesian learner is also sensitive to the discount rate. The primary difference between these learning models is that the Bayesian is able to eventually learn about the rate of climate change and therefore anticipate future change whereas the rolling-mean learner permanently lags the true climate state. Therefore, a lower (higher) discount rate will increase (decrease) the difference in $\tilde{C}_u$ between different learning models.

Figure 2(b) shows the spatial pattern of $\tilde{C}_{ul}$ for the two emissions scenarios and three learning cases. Spatial differences are due largely to differences in rates of warming (Fig. A.1(a)). The more rapid warming at higher latitudes and in continental interiors leads to larger adjustment costs. Spatial differences in natural variability,

(a)

Figure 2. (a) Global, population-weighted average learning-related adjustment costs ($\tilde{C}_u$) for 2010–2050 under two emissions scenarios and three learning models, using a discount rate of 2.5%. RCP 8.5 is a business-as-usual emissions scenario, whereas RCP 2.6 is an ambitious mitigation scenario. (b) Learning-related adjustment costs ($\tilde{C}_{ul}$, Eq. (1)) for 2010–2050 for two emissions scenarios and three learning models. Plotted is the average difference (in $^{\circ}\text{C}$) between expectations and the true mean climate state, discounted at a rate of 2.5%

Fig. 2. (Continued)

which is much higher in temperate compared with tropical regions (Fig. A.1(b)), also play a role for the Bayesian learner as higher natural variability makes the signal of a changing climate more difficult to identify and tends to increase adjustment costs. Figure A.2 compares the population-weighted mean values of $\tilde{C}_u$ for tropical and temperate regions. Costs are lower for all cases in tropics compared with temperate regions (Fig. A.2). This is due both to the somewhat slower rate of warming in the tropics and to the much lower inter-annual temperature variability. Costs for the nonlearner, which depend only on the rate of climate change, in tropical areas are 14.5–23.1% lower than those in temperate regions. For the Bayesian learner, costs depend both on the rate of climate change and on inter-annual variability

(noise), and the difference between temperate and tropical areas is twice as large (40.5–42.6%).

One other notable feature of this simulation is the absolute values of learning-related adjustment costs, which are remarkably small. The average discounted difference between beliefs and the true climate state for efficient Bayesian learners is less than one-tenth of a degree and is only 0.16°C for the rolling-mean learner under the business-as-usual scenario. Only nonlearners show what might be considered climatically meaningful differences of 0.45°C under the business-as-usual case or 0.23°C in the ambitious mitigation case. Whether or not these relatively small values result in economically meaningful adjustment costs depends on the specific set of adaptation strategies available to firms. Guo and Costello (2013) show that if adaptation technologies are discontinuous, then the value of adaptation is theoretically unbounded. This means that even small discrepancies between beliefs and the true climate could, in some circumstances, result in large adjustment costs. Hsiang (2016) in turn shows the converse: that if adaptation technologies are continuous, the envelope theorem implies that the effects of small differences in climate should be small. The small values found in these simulations may, therefore, suggest that learning-related adjustment costs are relatively small, although, as noted above, absolute values will be sensitive to the choice of discount rate.

4. Distinguishing Between Learning Models in a Laboratory Environment

The previous section has shown that how people learn from observations in a stochastic, nonstationary environment may be important for determining the adjustment costs associated with climate change.

However, empirically distinguishing between learning models is challenging because (1) in most areas the signal of anthropogenic climate change has yet to emerge from the noise of inter-annual variability (Hawkins and Sutton, 2012; Mora *et al.*, 2013) and (2) beliefs about the climate state can typically not be observed directly. An experimental setting, while abstracting away from some important aspects of the adaptation decision, can be useful because it provides a controlled environment that can be systematically manipulated in order to distinguish between learning models.

I use an online experimental game that simulates the decision to adapt to a changing climate by presenting subjects with a choice between two hypothetical crop varieties. Payoffs depend on the choice of crop and on the weather in a particular turn:

$$y_t = f(\theta_t^*, w_t),$$

where y_t are the points in turn t , $\theta_t^* \in [A, B]$ is the subject's ex-ante choice of Crop A or Crop B in turn t , and w_t is the random draw of weather for turn t . w_t is drawn from a normal distribution ("climate"), which is not observed by subjects and which changes over the length of the game, though with a fixed variance:

$$w_t \sim N(\mu_t, \sigma_w^2).$$

The two payoff functions ($f(\theta_A w_t)$) and ($f(\theta_B w_t)$) are both inverted quadratics, with the optimal weather for Crop B higher than that for Crop A. This set-up means that the subject faces the problem of choosing θ so as to maximize the expected number of points, conditional on subject's beliefs about the state of the climate:

$$\theta_t^* = \max_{\theta \in [A, B]} \int_{\underline{w}}^{\bar{w}} f(\theta, w_t) * P(w_t | \mathbf{w}_{t-1}) dw_t.$$

This means that observing the choice of crop (θ_t^*) over the course of the game allows inferences to be made about the evolution of the posterior predictive distribution ($P(w_t | \mathbf{w}_{t-1})$) and therefore what learning model is being used by the subjects.

I use two treatment groups to distinguish between a proactive learning model that anticipates future changes (such as the Bayesian model presented in this paper) and a reactive adaptation strategy that implicitly assumes that the climate is stationary (such as the rolling-mean model). Both groups face the same unobserved change in probabilities, an initial stationary baseline for 10 turns, followed by a linearly changing mean.

$$\mu_t = \begin{cases} \mu_0 & \text{for } t \leq 10, \\ \mu_0 + \beta(t - 10) & \text{for } t > 10. \end{cases}$$

I manipulate the payoff functions between treatment groups so that the optimal adaptation point varies. In both groups, payoff functions are such that the optimal choice under the baseline climate (μ_0) is Crop A. The optimal adaptation point (t^*) occurs when:

$$\int_{-\infty}^{\infty} f(\theta_A, w_t) * N(\mu_t, \sigma^2) dw_t = \int_{-\infty}^{\infty} f(\theta_B, w_t) * N(\mu_t, \sigma^2) dw_t.$$

In the “Early” treatment group, payoff functions are such that $t^* = 17$ (i.e., 7 turns after the “climate” starts changing) and in the “Late” treatment group $t^* = 29$.

A proactive learner is able to learn about a trend in the stochastic environment and so will be better able to anticipate the later change-point because they have had time to learn about the rate of change. This means that this kind of learner should adapt closer to the optimal adaptation point in the “Late” treatment group, relative to the “Early” treatment group. The beliefs of a learner using a reactive updating process such as the rolling mean do not converge to the true climate state. This means that this type of learner should show a similar delay in adapting between the two treatment groups. This difference is shown schematically in Fig. A.3. Figure A.4 shows the results of 100 simulations of the experiment under the hierarchical Bayesian model and the rolling-mean model, using parameters from the two experimental treatment groups. This shows large differences in adaptation time for the trend-learner, but no differences between treatment groups for the rolling-mean learner. A trend-learning model and the rolling-mean model can, therefore, be distinguished by testing whether the delay

between the optimal adaptation point (t^*) and the observed adaptation point (t_{obs}) is significantly less in the “Late” treatment group compared with the “Early” treatment group:

$$H_0: (t_{\text{obs}} - t^*)_{\text{Early}} = (t_{\text{obs}} - t^*)_{\text{Late}},$$

$$H_1: (t_{\text{obs}} - t^*)_{\text{Early}} > (t_{\text{obs}} - t^*)_{\text{Late}}.$$

A rejection of H_0 indicates that subjects faced with a gradually changing stochastic environment are able to learn about a trend in the environment and therefore to respond proactively to climate change.

The experiment was implemented using Amazon Mechanical Turk (AMT), with subjects drawn from the United States. Subjects were paid \$0.50 to participate, with much larger bonus amounts (of up to \$2.00) tied to performance in the game. All subjects played to completion, all took less than 10 min to complete the game, and each AMT worker only participated in the experiment once. Subjects were given an initial set of instructions on how to complete the game, including a histogram showing the baseline “climate” conditions. This was followed by 10 turns of the baseline conditions, followed by 40 turns of gradually changing conditions. The “Early” treatment group had 97 subjects and the “Late” group had 48. The observed adaptation point (t_{obs}) was determined by fitting a logistic curve to the fraction of participants choosing Crop B over the course of the game and observing the point at which this curve crossed 50%; 95% confidence intervals for $t_{\text{obs}} - t^*$ are derived from uncertainty in the parameters of these logistic curves.

Figure 3 shows the results of this experiment. In both treatment groups, subjects were more likely to choose Crop A at the beginning of the experiment and Crop B at the end of the experiment. This indicates that expectations about the “weather” distribution were being updated and therefore provides evidence against the no-learning model. Although subjects adapted in the game, indicating that learning was happening, adaptation was delayed in both treatment groups. The delay ($t_{\text{obs}} - t^*$) was 12.1 ± 2.4 turns in the early treatment group and 6.6 ± 1.9 turns in the late treatment group. A one-tailed hypothesis test indicates that the adaptation delay for the late treatment group is significantly shorter than that for the early treatment group ($p = 0.035$).⁵ This means that the null hypothesis can be rejected and that the results of this experiment provide some evidence that people are able to learn about a trend and anticipate future changes.

⁵The relevant hypotheses derived from the proactive versus reactive learning model, as well as the results of simulations shown in Fig. A.4, indicate that a one-tailed hypothesis test is appropriate here. However, previous papers that show comparable performance of simple learning rules relative to optimal strategies in some contexts (Klump, 2006) mean that it is possible that the hypothesis is mis-specified and a two-tailed test might be appropriate instead. A two-tailed test that the difference in delay is different between the early and late treatment groups rejects the null at the 10% level, but not at the 5% level ($p = 0.07$).

Figure 3. Difference between the observed adaptation point and the optimum adaptation point for two treatment groups in the experiment. Error bars show the 95% confidence interval

Because of the simplicity of the experimental set up, and the number of (potentially individual-specific) free parameters involved in fully specifying a learning model, this experiment is only able to distinguish between broad classes of models. Specifically, these results appear to rule out models in which subjects either do not react at all to a changing environment (i.e., no learning) or respond reactively to past changes but do not anticipate future changes (e.g., the rolling-mean model). The results suggest that subjects are able to recognize a trending climate and therefore respond proactively. This is consistent with the Bayesian model presented in this paper, but a number of other trend-learning models would also be consistent with the observed behavior.

5. Discussion and Conclusion

As the climatological distribution over weather shocks is not directly observed and because climate models can give us only a rough sense of how it is changing at a local scale, it seems likely that a primary source of information will be observations of weather, particularly unexpected (extreme) weather events. How rapidly people and firms adjust beliefs about the state of the climate limits the rate at which adaptation can occur and therefore is a first-order determinant of the adjustment costs associated with

climate change.⁶ Therefore, the accuracy and efficiency with which people use their own observations of weather to learn about a changing climate distribution will be an important determinant of total climate change damages, with significant implications for both adaptation and mitigation policy (Cropper and Oates, 1992; Stern, 2006).

This raises the question of how people update beliefs about an unobserved, non-stationary state variable, about which they receive only noisy information. This paper has presented a Bayesian model of learning in which observers learn both that climate change is happening and the rate at which it is changing, a more flexible formulation than the original learning model proposed in KKM. Comparing this with less computationally intensive learning model using parameters derived from climate models, I find substantial relative differences in the learning-related adjustment costs between learning models, although the absolute values of differences between the perceived and true climate distributions are small.

Whether these differences between the perceived and true climate distribution result in economically significant costs depends on the adaptation technologies available in individual regions and sectors. People prepare for the distribution of weather risks they believe they face, but if their beliefs are wrong, then they will over- or under-prepare for different weather events, both of which are costly relative to a situation with full information. If adaptation is discontinuous, then costs associated with imperfect adjustment are theoretically unbounded and could be substantial (Guo and Costello, 2013). However, if the difference between beliefs is small enough and the set of adaptation options discrete, then imperfect learning may result in no difference in decisions and therefore no welfare losses (Klumpp, 2006). In other cases in which adaptations are continuous, Hsiang (2016) points out that the envelope theorem implies that adjustment costs are unlikely to be economically meaningful, at least at the individual level.

The simple experimental result presented in this paper suggests that individuals may be able to learn about a gradually changing environment and respond proactively rather than simply react to changes that have already occurred. If people are able to anticipate future climate change and plan adaptations accordingly, then in some cases adjustment costs may be on the lower end of the possible range. However, many laws and regulations relevant to climate change adaptation implicitly incorporate a reactive, rolling-mean learning model. For example, the yield basis for crop insurance in the United States is typically calculated using a 10-year trailing mean (Plastina and Edwards), or the definition of a floodplain for zoning purposes might be based on the

⁶It may be that this learning process is also important in determining public support for mitigation policies, although the direction of this relationship is ambiguous. If observers realize that their expectations have adjusted, and attribute this change to anthropogenic climate change, this learning process may increase support for mitigation. If this learning process is unconscious though, so that previously unusual or extreme weather comes to be seen as normal, it may be that rapid adjustment will decrease support for mitigation. The question of how people attribute observed trends is beyond the scope of this paper. It is not necessary that trends be attributed to anthropogenic climate change for adaptation to occur, but it is an important question for mitigation.

historical 50- or 100-year flood events. Under a changing climate, measures of weather risk based on these metrics will always be biased. Therefore, this work points to the importance of having laws and institutions that learn about and anticipate a changing climate in order to promote efficient adaptation and minimize adjustment costs.

A range of recent works in the climate sciences has begun connecting the local weather events that people actually experience with the more abstract and intangible concept of global climate change. Relevant examples include literature on the rate at which the climate change signal emerges from natural variability (Diffenbaugh and Scherer, 2011; Hawkins and Sutton, 2012; Mora *et al.*, 2013), the attribution of local extremes to global climate change (Fischer and Knutti, 2015; Horton *et al.*, 2015), as well as work more explicitly relating people's experience of extremes with perception of global climate change (Ricke and Caldeira, 2014; Lehner and Stocker, 2015). In understanding what experience of weather actually means for people affected, it is important to remember that people can, and almost certainly will, adjust their expectations of and preparations for weather as the climate changes: what once seemed unusual may eventually be considered normal, and once normal weather may eventually become extreme. Although some kind of learning will almost certainly take place, the exact manner of learning and the associated rate of adaptation are still poorly understood. This paper then points to the importance of future theoretical and, in particular, empirical extensions in order to understand the formation of expectations in a changing climate and to better constrain adaptation costs.

Acknowledgments

The author thanks the Neukermans Family Foundation for support of this research. The author also thanks David Lobell and participants in the Stanford Environmental and Energy Policy Analysis Center seminar for comments and suggestions.

References

- Akerlof, G and J Yellen (1985). Can small deviations from rationality make significant differences to economic equilibria? *American Economic Review*, 1(75), 708–720.
- Adams, RP and DJC MacKay (2007). Bayesian online changepoint detection. *arXiv*, 1050(1), 7.
- Burke, MB and K Emerick (2016). Adaptation to climate change: Evidence from U.S. agriculture. *American Economic Journal: Economic Policy*, 8(3), 108–140. Retrieved from <http://web.stanford.edu/~mburke/papers/Burke.Emerick.2015.pdf>.
- Cropper, ML and WE Oates (1992). Environmental economics: A survey. *Journal of Economic Literature*, 30(2), 675–740.
- Diffenbaugh, NS and M Scherer (2011). Observational and model evidence of global emergence of permanent, unprecedented heat in the 20th and 21st centuries. *Climatic Change*, 107(3–4), 615–624. <http://doi.org/10.1007/s10584-011-0112-y>.

- Fischer, EM and R Knutti (2015). Anthropogenic contribution to global occurrence of heavy-precipitation and high-temperature extremes. *Nat Climate Change*, 5(6), 560–564. <http://doi.org/10.1038/nclimate2617>.
- Gilovich, T, R Vallone and A Tversky (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17, 295–314.
- Guo, C and C Costello (2013). The value of adaption: Climate change and timberland management. *Journal of Environmental Economics and Management*, 65(3), 452–468. <http://doi.org/10.1016/j.jeem.2012.12.003>.
- Hawkins, E and R Sutton (2012). Time of emergence of climate signals. *Geophysical Research Letters*, 39(1), L01702. <http://doi.org/10.1029/2011GL050087>.
- Hornbeck, R (2012). The enduring impact of the American Dust Bowl: Short and long-run adjustments to environmental catastrophe. *American Economic Review*, 102(4), 1477–1507.
- Horton, DE, NC Johnson, D Singh, DL Swain, B Rajaratnam and NS Diffenbaugh (2015). Contribution of changes in atmospheric circulation patterns to extreme temperature trends. *Nature*, 522(7557), 465–469. <http://doi.org/10.1038/nature14550>.
- Howe, PD, H Boudet, A Leiserowitz and EW Maibach (2014). Mapping the shadow of experience of extreme weather events. *Climatic Change*, 127(2), 381–389. <http://doi.org/10.1007/s10584-014-1253-6>.
- Hsiang, SM (2016). Climate econometrics. Retrieved from <http://www.nber.org/papers/w22181>.
- Kahneman, D (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, 93(5), 1449–1475.
- Kahneman, D, JL Knetsch and RH Thaler (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. 5(1), 193–206.
- Kala, N (2016). Learning, adaptation and climate uncertainty: Evidence from Indian agriculture. Retrieved from https://namratakala.files.wordpress.com/2014/01/kala_learning-jan20171.pdf.
- Kaufmann, RK, ML Mann, S Gopal, JA Liederman, PD Howe, F Pretis, ... M Gilmore (2017). Spatial heterogeneity of climate change as an experiential basis for skepticism. *Proceedings of the National Academy of Sciences of the United States of America*, 114(1), 67–71. <http://doi.org/10.1073/pnas.1607032113>.
- Kelly, DL and CD Kolstad (1999). Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control*, 23(4), 491–518.
- Kelly, D, C Kolstad and G Mitchell (2005). Adjustment costs from environmental change. *Journal of Environmental Economics and Management*, 50(3), 468–495. <http://doi.org/10.1016/j.jeem.2005.02.003>.
- Klump, T (2006). Linear learning in changing environments. *Journal of Economic Dynamics and Control*, 30, 2577–2611.
- Lehner, F and TF Stocker (2015). From local perception to global perspective. *Nat Climate Change*, 5(August), 731–734. <http://doi.org/10.1038/nclimate2660>.
- Marvel, K and C Bonfils (2013). Identifying external influences on global precipitation. *Proceedings of the National Academy of Sciences*, 110(48), 19301–19306. <http://doi.org/10.1073/pnas.1314382110>.
- McCright, AM, RE Dunlap and C Xiao (2014). The impacts of temperature anomalies and political orientation on perceived winter warming. *Nat Climate Change*, 4(12), 1077–1081. <http://doi.org/10.1038/nclimate2443>.
- Mendelsohn, R (2000). Efficient adaptation to climate change. *Climatic Change*, 45, 583–600.

- Mendelsohn, R, WD Nordhaus and D Shaw (1994). The impact of global warming on agriculture: A Ricardian analysis. *American Economic Review*, 84(4), 753–771.
- Mora, C, AG Frazier, RJ Longman, RS Dacks, MM Walton, EJ Tong, ... TW Giambelluca (2013). The projected timing of climate departure from recent variability. *Nature*, 502(7470), 183–187. <http://doi.org/10.1038/nature12540>.
- Plastina, A and W Edwards (n.d.). Proven yields and insurance units for crop insurance. Retrieved July 22, 2015, from <http://www.extension.iastate.edu/agdm/crops/html/a1-55.html>.
- Rhines, A and P Huybers (2013). Frequent summer temperature extremes reflect changes in the mean, not the variance. *Proceedings of the National Academy of Sciences*, 110(7), E547–E548.
- Ricke, KL and K Caldeira (2014). Natural climate variability and future climate policy. *Nature Climate Change*, 4(5), 333–338. <http://doi.org/10.1038/nclimate2186>.
- Simon, HA (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Stern, N (2006). *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press.
- Taraz, V (in press). Adaptation to climate change: Historical evidence from the Indian monsoon. *Environment and Development Economics*.
- Thompson, DWJ, EA Barnes, C Deser, W Foust and AS Philips (2015). Quantifying the role of internal climate variability in future climate trends. *Journal of Climate*. <http://doi.org/10.1175/JCLI-D-14-00830.1>.
- Tversky, A and D Kahneman (1974). Judgment and uncertainty: Heuristics and biases. *Science*, 185, 1124–1131.
- Wilson, PS and R Toumi (2005). A fundamental probability distribution for heavy rainfall. *Geophysical Research Letters*, 32(14). <http://doi.org/10.1029/2005GL022465>.