COLUMBIA | SIPA Center for Environmental Economics and Policy

CEEP WORKING PAPER SERIES WORKING PAPER NUMBER 17

APRIL 2021

Adjusting to Rain Before It Falls

Mitch Downey, Nelson Lind, and Jeffrey G. Shrader

https://ceep.columbia.edu/sites/default/files/content/papers/n17.pdf

Adjusting to Rain Before It Falls

Mitch Downey, Nelson Lind, and Jeffrey G. Shrader^{*}

April 22, 2021

Abstract

Unchecked climate change will cause precipitation volatility to increase around the world, leading to economic damages in the face of adjustment costs. We estimate these damages for construction, an economically important, climate exposed industry. Empirically, employment falls in response to forecasted rainfall and more so as the forecast horizon increases. This pattern allows us to identify labor adjustment costs. Calibrating a multi-sector model of local labor markets to match our estimates, we find that a typical rainfall event leads to an 11% loss in construction sector productivity, which firms anticipate and adapt to through costly adjustment of their labor force. Higher adjustment costs reduce this adaptation, leading to greater damage from bad weather realizations. These results imply that firms value forecasts and would be willing to pay 3.6 months of typical profits per year to learn about rainfall 5 months sooner. (JEL:D83,J21,Q51)

^{*}Downey: Institute for International Economic Studies, Stockholm University (email: mitch.downey@ iies.su.se); Lind: Department of Economics, Emory University (email: nlind@emory.edu); Shrader: School of International and Public Affairs, Columbia University (email: jgs2103@columbia.edu). We are grateful for comments from seminar and conference participants at AERE Summer Conference, Cornell University, Emory University, the Philadelphia Federal Reserve Bank, UCLA, University of Southern California, and University of Virginia. Gabriel Gonzalez Sutil and Zeyba Qamar provided excellent research assistance. All errors are our own.

1 Introduction

Beyond simply warming the earth, climate change will have profound effects on the complete distribution of weather. While much of the literature has examined the economic consequences of changes in average weather (including rising temperatures and sea levels), in this paper we consider economic effects of changes in the *volatility* of weather. In particular, we focus on productivity interruptions caused by rainfall, show that firms incur serious costs from the types of rainfall shocks that climate change makes more common, and study the extent to which scientific advances in forecasting and attentive planning by managers can offset these costs.

We focus on the construction sector, a large and economically important industry that is highly exposed to the climate. Construction constitutes a five times larger share of GDP in the U.S. than the more widely studied forestry, fishing, and agriculture sector (BEA, 2019).¹ Beyond simply its size, construction is central in input-output networks, supplying essential inputs to virtually all other industries and using materials produced by many other sectors. At the same time, weather directly impacts construction. Data from the American Time Use Survey shows that only agricultural workers spend more time outside than construction workers. Combined with the importance of construction to the overall economy, month-tomonth disruptions due to weather can be an important source of economic costs.

Despite its economic importance and potential climate vulnerability, construction has not been the focus of research in the climate economics literature.² This literature has focused either on the costs associated with equilibrium consequences of gradual changes in average temperature and precipitation or on the acute effects of realized weather shocks, both of which have first order effects on sectors like agriculture and energy.³ In contrast, we focus on the costs regularly paid due to *ex ante* adjustment to weather events, and we use evidence from the construction sector to infer the increase in these costs that will occur after climate change has shifted weather patterns.

Specifically, climate change will increase rainfall volatility as warmer air masses hold and then suddenly release larger volumes of water (Pendergrass et al., 2017). This effect

¹This multiple is not an outlier. According to UNECE (2019), the construction industry is 9 times larger than agriculture in the U.K. and 6 times larger in Germany. Even in more heavily agriculturally reliant countries like Spain and Italy, construction is more than twice the size of agriculture in terms of its contribution to GDP.

²For the few studies that do look at construction, the emphasis is often on potential job growth as other sectors comply with climate policy or adapt to a changing climate. See Fankhaeser et al. (2008), for example. Some research studies sectoral effects of weather and climate change. Graff Zivin and Neidell (2014) estimate the effects of temperature on intensive-margin labor supply in construction and other climate-exposed industries. Jain et al. (2020) investigate the effects of temperature on economic production for multiple sectors in India including construction.

³Zhang et al. (2017) points out that a wide range of climate variables—not just temperature and precipitation—will likely have important effects on agriculture.

is often summarized by the saying that climate change will cause "wet places to get wetter and dry places to get drier." Climate projections also show that within locations, the time series volatility will increase (O'Neill et al., 2016).⁴ Consistent with a large body of evidence on the costs of volatility in conditions like product demand or supply chains, we show that productivity volatility caused by rainfall shocks is an important but neglected cost of climate change.⁵

We empirically identify adjustment costs in the construction industry by estimating employment responses to more or less forecastable rainfall. Our primary focus is on how medium-range forecasts of weather conditions in the coming six months affect the dynamics of employment responses to rainfall. We combine data on construction employment in commuting zones across the United States over the last three decades with: i) information on monthly rainfall and ii) climate variation that can be used to forecast rainfall month-tomonth. These data allow us to estimate employment responses to exogenous news shocks about future rainfall. Our identification strategy exploits rich variation in weather patterns across the U.S. to identify separate employment responses to forecasts available at different horizons (ranging from one month to six months in advance). To help validate our empirical approach, we provide extensive discussion of the ways in which rainfall is a serious and wellrecognized challenge in construction sector, construction firms actively monitor the types of medium-range forecasts we rely on for identification, and commonly used contracts shift weather risk onto construction firms.

We find that the elasticity of construction employment with respect to rainfall is five times larger when that rainfall could be predicted 6 months in advance compared to when it could only be predicted 1 month in advance. In other words, an identical rainfall "shock" generates a significantly larger employment response when it is predictable further in advance. This empirical reduced form finding cannot be reconciled with simple, frictionless, static models of labor input, which is important partly because these models remain the norm for analyzing economic effects of climate change.

To interpret our reduced form effects, we model a multi-sector, open-economy with both demand-side and supply-side adjustment costs. We calibrate our model to match the employment dynamics that we estimate. Unlike the benchmark models widely used for climate economics, our model creates a channel by which volatile productivity can harm firms (because they wish to adjust their input demands but must face costs of doing so). We use this model to understand how productivity and output are affected by rainfall shocks, how the

 $^{^{4}}$ Over the last 30 years, month-to-month variance in rainfall has increased by about 5% in the U.S. CMIP6 projections indicate that variance will go up another 10 to 15% by the end of the century under unmitigated warming.

⁵Fishman (2016), in a rare evaluation of the effects of rainfall variability, finds that climate-change-driven increases in rainfall volatility (leading to a higher probability of dry days) will offset the benefit of increasing total rainfall.

burden of these shocks is shared between workers and firms, and the extent to which firms are able to adjust employment to avoid the costs of the shock. Throughout the counterfactuals that we present, we pay special attention to how our results would change if firm-side adjustment costs fall (for instance, because firms write more flexible employment contracts that allow greater responsiveness to shocks) or the forecast horizon for the rainfall fluctuation increases (for instance, because managers increasingly use better scientific models).

In our benchmark calibration, we find that rainfall shocks reduce firm profits. However, this effect varies by forecast horizon. A firm that is able to forecast the shock six months in advance can offset 80% of the profit loss that would occur if the it was only able to forecast the shock one month in advance. Moreover, if firms face sufficiently low adjustment costs, then they can actually *benefit from* the rainfall shock because the transitory productivity decline leads to a pent-up demand boom after the rain has passed. At the same time, these features create a real policy trade-off. Firms largely offset the effects of the productivity shock by passing the costs onto workers. Therefore, just as longer forecast horizons and more flexible adjustment make firms better off, they leave workers worse off.

The firms do pay a cost to engage in this adaptation, however. There is a direct cost stemming from changes to the labor force and the adjustment costs we estimate. These costs are always weakly positive, and they occur whether the rainfall shocks help or hurt productivity. The level of adjustment cost also determines how willing the firm is to engage in *ex ante* labor adjustment in order to reduce damages when the shock arrives. In our baseline calibration, we show that adjustment costs are high enough to lead to non-trivial adaptation costs for firms, but low enough to make firms want to engage in substantial adjustment—particularly if they have enough advance notice before the shock arrives. The net benefit of forecasts is large enough that over the course of a year, firms would be willing to pay more than three months of typical profits to replace 1-month ahead forecasts with 6-month ahead forecasts.

If rainfall gets more volatile in the future, as climate projections currently indicate, then the adjustment costs we estimate will translate into extra damage for the economy. Conditional on a rainfall shock occurring, higher volatility means that a shock will have a higher probability of being larger in magnitude. In the absence of adaptation, larger shocks will cause bigger losses in productivity and profit for firms. Our estimates, however, show that firms can offset some of those losses by planning further ahead. Given the projected 10 to 15% increase in rainfall volatility by the end of the century, a firm would require up to a half-month further ahead forecast to be left no worse off. This might seem like a short increase in horizon, but due to non-linearity in weather systems, the difficulty of skillfully forecasting at longer horizons rises exponentially.

Thus, these results might serve as a warning about future climate damages while also

pointing to two avenues firms can explore to reduce their risks. Construction managers know that their bottom line is affected by weather and have substantial experience planning for adverse weather (Trauner et al., 2018). If managers can further improve the quality and accuracy of their plans, then they can reduce the losses they will take if bad weather arrives. In-house project managers, weather risk consulting firms, and professional forecasters might all be able to aid in this effort. Improvements in forecasting at the horizons we study here, however, are challenging (Toth and Buizza, 2019). And as the results on workers show, forecast improvements come with the potential for increasing inequity in the incidence of climate damages. Second, managers can potentially invest in ways that will reduce the adjustment costs they face. The construction industry already routinely uses seasonal contracts to manage the workforce during the winter (Krane and Wascher, 1999, Organization, 2016). A contract for shorter-run weather shocks that resembles these seasonal contracts could, for instance, reduce the costs the firm needs to pay to build back up the labor force after the shock dissipates.

Methodologically, our paper provides a unique empirical strategy to identify labor adjustment costs and to draw the link between those costs and climate change. There is a nascent literature focusing on the effects of climate change on labor markets, primarily in agriculture (for instance, see Rosenzweig and Udry (2014), Colmer (2020)). The current paper shows that climate change is important for extensive margin labor adjustment in construction as well. Effects in construction might be particularly relevant because construction services—like seawalls, retrofits of buildings to improve HVAC systems, accommodation of changing patterns of urbanization—will be important components of adaptation if climate change continues (Fried, 2019).

At the same time, our results show that adjustment costs in the construction sector mean that an important part of the effect of climate change on the industry happen before weather realizations occur. Failing to take this forward-looking component into account can bias estimates of the effect of climate change (Shrader, 2020, Lemoine, 2018). The methodology from this paper is directly portable to other locations around the world where monthly or seasonal forecasts are available (organizations like NOAA and the ECMWF now routinely provide such forecasts globally (Scher and Messori, 2019)).

Finally, our paper investigates a novel dimension of the damage from climate change. In this paper, when we say "adjustment costs," we are referring to costs routinely paid, day to day, that arise because of a widening of the weather distribution. Importantly, these are not adjustment costs paid when transitioning between equilibria. Instead, they are costs that are already being paid in steady state, every time rain arrives. Under projected climate change, we will simply be paying them more often. These costs are missing from existing analyses of the damages of climate change. Related to this point, as an added empirical contribution of our paper, the short-run nature of these costs mean that existing data speaks more directly to them than to costs arising from differences in long-run equilibria.

2 Related literature and background

2.1 Previous research

Volatility is central to business decisions. Volatility in either productivity or product demand has been shown to explain differences in investment behavior (Kellogg, 2014), market size and competition (Collard-Wexler, 2013), employment contracts and outsourcing decisions (Abraham and Taylor, 1996), and the "misallocation" of factors of production (Asker, Collard-Wexler, and De Loecker, 2014).⁶

However, beyond its direct implications for firms' operational decisions, the importance of volatility is informative about the fundamental structure of the labor market. Models with a frictionless, perfectly competitive labor market (still the foundations of many policy analyses, including climate models) leave no room for volatility. If firms can freely adjust their labor force, as these models assume, then employment and profits depend only on the actual conditions firms are facing at any given moment. Changes in those conditions, regardless of their size, frequency, or predictability, are not important because firms can costlessly and immediately re-optimize.

It has long been understood that this is not realistic. Firms face dozens of complex costs in changing their labor pool.⁷ The private sector has understood this longer than academic economists have. In his seminal 1962 paper, Walter Oi (1962) argues that non-negligible hiring and training costs mean that part of a firm's labor is a fixed cost. Oi rests this argument largely on an internal study done a decade earlier by the International Harvester Company (IHC) called "The Costs of Labor Turnover" (1951). There, IHC economists estimate that the average cost of training a new worker was \$238, or about 11% of median annual earnings at the time (\$2,200). These costs have remained remarkably persistent. Representative establishment surveys conducted today estimate that the costs of replacing an employee are 9% of average annual earnings (Dube, Freeman, and Reich, 2010).

⁶The bulk of empirical studies of productivity, including those focused on volatility, treat it as a residual (i.e., excess production after accounting for inputs). The estimation challenges in this approach are substantial and well-understood (Ackerberg, Caves, and Frazer, 2015). A contribution of our paper, made possible by our context, is that we study productivity shocks using a direct measure of productivity: rainfall, which has first order effects on the ability to perform construction, as discussed below.

⁷Our focus is on the effects of increasing the *volatility of* productivity fluctuations, since this is directly linked to climate change. A separate but related issue is the effects of *uncertainty about* future productivity fluctuations. As Bloom (2009) points out, uncertainty about future productivity also has no effect in a frictionless model, but in a model with realistically calibrated adjustment costs, can have first order effects on aggregate employment and output.

Of course, the costs facing a firm wishing to adjust its labor pool extend beyond recruiting and training workers. A large literature on wrongful discharge regulations shows that these regulations impose firing costs. By making it more difficult for firms to dismiss workers when they need to, firing costs have been shown to reduce firms' hiring, turnover, ability to respond to shocks, and ultimately their financial value (see Serfling (2016) for a review).⁸ Independent of these regulations, all US firms' face direct financial costs of layoffs arising from the financing rules of the Unemployment Insurance (UI) system (Ratner, 2013).⁹

More generally, a range of diverse evidence—including the costs firms incur when they suddenly and unexpectedly lose an employee (Ginja, Karimi, and Xiao, 2020, Isen, 2013, Jäger and Heining, 2019), the value to firms of hiring through employee referral networks (Burks et al., 2015), and employers' learning about employee skills over time (Kahn and Lange, 2014)—suggests that firms cannot easily shrink or grow their workforce without costs: a firm's current workers and the potential alternatives "out there" are imperfect substitutes.¹⁰ At a minimum, given incomplete markets and imperfect insurance for risk, a firm that is likely to suddenly cut its workforce must still compensate its workers (through higher wages) for accepting this risk, an idea which dates back to Adam Smith (1776).¹¹

Against this backdrop, our core contributions are, first, to estimate the magnitude of adjustment costs in the construction sector using a novel strategy that exploits variation in whether and when rainfall-driven productivity fluctuations could have been predicted, and second, to quantify the consequences of increasing rainfall volatility driven by climate change. An advantage of our method for estimating adjustment costs is that it comprehensively includes all types and sources of adjustment costs. Analyses of specific laws (like wrongful

⁸A paper closely related to ours is Adhvaryu, Chari, and Sharma (2013), who study how agricultural employment responds to rainfall in India. They focus on differential responses depending on wrongful discharge laws and their implications for labor regulation. We focus on differential responses depending on whether and when that rainfall could have been predicted and the implications for rising volatility driven by climate change.

⁹While much of the literature emphasizes regulations as the source of frictions and adjustment costs, 19th century manufacturing wages show evidence of imperfect competition in the labor market and firm-specific rents, even before labor regulations existed in the U.S. (Naidu and Yuchtman, 2018).

¹⁰It has long been understood that one challenge facing firms is screening applicants for quality. There is growing evidence that non-cognitive skills and traits are becoming increasingly important in the labor market (Edin et al., 2017), and this matters because it is likely even more difficult to screen for these skills than for cognitive or training-based skills. Many of these non-cognitive traits are directly relevant for construction firms. For instance, construction firms need workers who are reliable and will arrive on time, who can get along with others in an inherently interactive job, and who respect and value safety practices, all of which are extremely difficult to assess during the hiring process. This increases the incentive for construction firms to keep their workers and the challenges of changing staffing levels in response to demand fluctuations.

¹¹ "In the greater part of manufactures, a journeyman may be pretty sure of employment almost every day of the year that he is able to work. A mason or bricklayer, on the contrary, can work neither in hard frost nor in foul weather, and his employment at all other times depends upon the occasional calls of his customers. He is liable, in consequence, to be frequently without any [employment]... The high wages of those workmen, therefore, are not so much the recompence of their skill, as the compensation for the inconstancy of their employment."

discharge laws (Serfling, 2016) or UI financing (Ratner, 2013)) only capture one specific source of adjustment costs, and survey-based approaches (like Oi (1962) and Dube et al. (2010)) can easily miss important types of costs. An advantage of our application to climate change is that our estimation strategy is explicitly and directly linked to our counterfactual of interest (month-to-month volatility of rainfall), and outlines clear implications for firms and policy to respond (improving the ability to forecast future rainfall).

2.2 Industry attention to rainfall and long-range forecasts

The construction industry is a natural and important setting to estimate the effect of climate volatility—particularly from rain—on employment and productivity. Construction firms face strong incentives to pay attention to and plan for rain. Rain delays the completion of construction projects, and the costs of those delays typically fall on the construction contractor.

Rainfall makes outdoor work more hazardous by reducing visibility and increasing the risk of slipping or falling. Rain generates mud that can impede access to a work site and prevent the use of heavy machinery. It also prevents certain types of welding, electrical work, and cement pouring. Rain can even delay indoor work. For example, heavy rainfall can cause the water table to rise, delaying basement construction. Finally, construction firms in the United States are required by the Environmental Protection Agency to control rain-driven pollution effluent from work sites, requiring the diversion of labor and capital to that task (Environmental Protection Agency, 2009). One recent estimate by Ballesteros-Pérez et al. (2018) shows that weather variation delays the average construction project by 22%.

Under common contracts, construction firms bear much of the risk for weather-related delays (Trauner et al., 2018). The American Institute of Architects (AIA) Form A201—a standard contract between a construction contractor and a site owner that is widely used in the U.S.—stipulates that the contractor is responsible for any rain delay that could have been "reasonably anticipated."¹² Even if a delay is caused by potentially unreasonable weather, the contractor must still bring a claim and show, as a matter of fact,¹³ that the weather was more extreme than could have been expected. This finding of fact requires that the contractor record and pay attention to weather conditions during the construction project.¹⁴ Even if a construction contractor successfully makes a claim for an extension to a

¹²In the contract, construction firms can attempt to lengthen the contract without incurring cost under A201 §15.1.6.2, which reads: "If adverse weather conditions are the basis for a Claim for additional time, such Claim shall be documented by data substantiating that weather conditions were abnormal for the period of time, could not have been reasonably anticipated, and had an adverse effect on the scheduled construction."

 $^{^{13}}A$ matter of fact is a legal term for an actual occurrence, in contrast to a matter of law which is the purview of the judge.

¹⁴If not otherwise stated by contract, the presumption in the courts is that weather risk falls squarely on the contractor. Associated Engineers and Contractors v. State, 58 Haw. 322: "Moreover, the risk of

delayed project, the contractor might simply avoid paying for the delay without gaining any additional recompense. Moreover, beyond the within-project costs, delays in one project can have knock-on effects of other projects that make them even more costly.

Given the high cost of weather delays and the contractual need to form "reasonable" expectations about those delays, construction firms devote substantial resources to planning for weather. Bids at the beginning of jobs routinely include a monthly breakdown of expected weather delay days, and the construction project manager is expected to record actual delays relative to this schedule (Trauner et al., 2018).¹⁵ Specialized, proprietary project management software helps review weather data before and during a project.¹⁶

The key implication of all this is that the construction industry pays particular attention to *long-range* weather forecasts and the El Niño/Southern Oscillation (ENSO) variation that we use in our analysis is the most important element of those. ENSO is a coupled oceanicatmospheric phenomenon that occurs in the equatorial Pacific Ocean and is a primary driver of medium term global climate variation (Ropelewski and Halpert, 1987). El Niño events lead to higher rainfall in most (but not all) of the U.S.¹⁷ ENSO events change weather sometimes for months thereafter, allowing forecasters to make skillful predictions of weather patterns over monthly, seasonal, and annual horizons. The European Centre for Medium-Range Weather Forecasts (ECMWF), the premier weather forecasting group in the world, states that "Long term predictions [extending out to seven months] rely on aspects of Earth system variability which have long time scales (months to years) and are, to a certain extent, predictable. *The most important of these is the ENSO (El Niño Southern Oscillation) cycle.*" (emphasis ours, ECMWF (2019)). NOAA also issues seasonal weather forecasts for the U.S., again relying heavily on predictability coming from ENSO.¹⁸

A strong El Niño event in 1982 and 1983 caused substantial rainfall across the U.S. and, over the next decade, led to increased public awareness of the link between ENSO and weather conditions.¹⁹ The construction industry is well aware of the link between ENSO and rainfall. Articles in construction trade journals routinely report on the link between ENSO

abnormal weather is commonly held to be assumed by a construction contractor, except where provision otherwise is made in the contract or the parties are not equal in their knowledge of relevant weather data. Hardeman-Monier-Hutcherson v. United States, 458 F.2d 1364 (Ct. Cl. 1972)."

¹⁵These project managers make up 9% of all employment in the construction industry and hold college degrees at similar rates to the full US workforce, with engineering degrees being among the most common.

¹⁶Some contracts, such as the State of Tennessee RPA January 2002 Std 01252, specifically requires that the contractor consult National Oceanic and Atmospheric Administration data to determine anticipatable weather delays.

¹⁷ENSO also causes changes in U.S. temperatures, but this effect is generally weaker than the precipitation effect (Ropelewski and Halpert, 1987, Halpert and Ropelewski, 1992).

¹⁸For the latest seasonal forecasts from NOAA, extending out to 12.5 months, see here: https://www.cpc.ncep.noaa.gov/products/predictions/90day/

¹⁹See, for instance, the Washington Post's history of ENSO reporting by various new agencies (Williams, 2015).

and U.S. rainfall, including emphasizing the need to prepare for changes in rainfall due to the climactic variation (Halsey, 2016). In summary, then, the ENSO fluctuations that we use to generate variation in predictable rainfall are both the major driver of long-run rainfall forecasts *and* are a major focus of construction sector project managers.

3 Data

To estimate the effect of more or less forecastable rainfall on construction employment, we combine three primary data sources on employment, weather, and climate variability.

3.1 Construction employment data

We measure employment using the Quarterly Census of Employment and Wages (QCEW), which provides a high frequency snapshot of employment across U.S. counties based on state Unemployment Insurance records. Although we are primarily interested in firms' responses, we focus on aggregate data at the labor market level because existing firm-level datasets are not suitable for our purposes. Nearly all firm-level datasets are annual, while our approach requires high-frequency employment adjustments. Fortunately, the QCEW provides employment at the monthly frequency.

While the Longitudinal Employer-Household Dynamics (LEHD) includes *quarterly* employment information (and would have other advantages), only since 2003 has it included all states. Prior to that, it includes only some states. This is a problem for us because part of our identifying variation comes from geographic variation in responsiveness to ENSO (thus, we need data from all states to provide enough cross-sectional variation) and part of our identifying variation comes from over-time variation in the ENSO index (thus, we need more over-time variation than the post-2003 period could provide). Figure A2 shows ENSO anomalies since 1990. Restricting to the post-2003 period would cost us the two highest ENSO spikes and the two lowest ENSO troughs.

3.2 Weather and climate data

The second main dataset we use contains rainfall and temperature measurements from the PRISM (Parameter-elevation Regressions on Independent Slopes) Climate Group (PRISM Climate Group, 2004). PRISM combines weather station data with elevation data to produce monthly, gridded measures of weather. The PRISM data provide more consistent geographic coverage than raw weather station data.

We aggregate the gridded measures to the commuting zone level by calculating populationweighted averages. The population weights come from the 2010 U.S. Census population grid available from the Center for International Earth Science Information Network (CIESIN, 2017). Figures showing the spatial variation in weather can be found in the appendix (Section A). The final weather dataset is a monthly panel from January 1990 to December 2016 of population-weighted average values for each commuting zone in the continental U.S. for total precipitation and average temperature.

Third, we use monthly data on El Niño/Southern Oscillation (ENSO) from the National Oceanic and Atmospheric Administration (NOAA) as a source of long-range (monthly) predictability in rainfall. ENSO is commonly measured using sea surface temperature anomalies in an area of the equatorial Pacific Ocean known as the Niño 3.4 region that extends from 5°S to 5°N latitude and 170°W to 120°W longitude. Warm anomalies in this region are classified as El Niño events and cold anomalies are classified as La Niña events. We use the Niño 3.4 index as our measure of ENSO in the paper. The history of the Niño 3.4 index over our sample period is shown in Figure A2 (it is worth noting that ENSO is acyclical with respect to U.S. recessions).

As discussed in Section 2.2, ENSO is a crucial component of monthly or seasonal forecasts released by NOAA, ECMWF, and other forecasting groups. Technological advances concentrated during the 1980's led NOAA to begin releasing routine forecasts of monthly weather conditions starting in the middle of 1989 (Shrader, 2020). The timing of release of these forecasts and the growing public attention to ENSO-driven weather in the U.S. motivates our focus on the period after 1989.

3.3 Estimation sample and summary statistics

Combining all of the above datasets, we have an estimation sample consisting of monthly, commuting zone-level observations of employment in the construction industry, rainfall, temperature, and the Niño 3.4 index. The dataset runs from January 1990 through December 2016. After excluding commuting zones with suppressed construction employment, the final sample includes 633 out of the 722 commuting zones in the continental U.S.

Table 1 shows summary statistics for the main variables in our analysis. Employment is our primary outcome of interest. Rainfall, as forecasted by ENSO, is our primary right-hand side variable. Temperature is a control in the primary analysis. With 633 disclosed CZs observed monthly for 27 years, the final sample consists of 205,092 total observations.

Statistics on excluded commuting zones are also reported in panel (b) of the table. The first row shows that the CZs in our sample account for 93% of all reported private sector employment in the QCEW (which is subject to minimal suppression). The disclosed CZs account 99% of total population. These two figures shows that the suppressed CZs are those that have minimal population and relatively small levels of employment. The final row shows that the disclosed sample is about 88% of all continental U.S. CZs. Figure A1 shows a map of the non-disclosed CZs. One can see that they are generally sparsely populated locations

(a) Estimation variables								
Variable	Mean	Std. Dev.	Obs.					
Construction employment	8604.4	22504.6	205,092					
Monthly rain (mm)	78.4	61.9	$205,\!092$					
Monthly avg. temperature ($^{\circ}C$)	12.7	9.9	$205,\!092$					
Niño 3.4 index (°C)	0.045	0.89	$205,\!092$					
	~	1 05						

Table 1: Summary Statistics

(b) Disclosed vs. Suppressed CZs							
Disclosed Suppressed Pct. disclosed							
Number of commuting zones	633	90	88%				
Total private employment (millions)	185.8	15.0	93%				
Year 2000 Population (millions) 276.3 2.7 99%							

(a) Estimation Variables

Notes: The table shows summary statistics for the estimation sample (panel a) and information on the disclosed CZs that are included in the estimation sample versus the non-disclosed CZs that are excluded (panel b). The estimation sample consists of a balanced panel of 633 CZs observed for 324 months, resulting in 205,092 total observations.

in the West.

4 Results

4.1 Predicting rainfall

Our core empirical specifications are instrumental variables (IV) regressions. The first stage in each specification generates predictions of rainfall using variation in ENSO. The first stage for a given forecasting horizon ℓ is a regression of the form

$$\ln(Rain)_{c,t} = \delta_t^1 + \rho_{c,m(t)}^1 + \gamma_c ENSO_{t-\ell} + \eta_c^1 ENSO_{t-\ell-6}$$
(1)
+ $\theta_{c,1}^1 \ln(Rain)_{c,t-1} + \theta_{c,2}^1 \ln(Rain)_{c,t-2} + \sum_{k=3}^{12} \pi_k^1 \ln(Rain)_{c,t-k}$
+ $\sum_{k=1}^{12} \zeta_k^1 Temp_{c,t-k} + \nu_{c,t}$

where c indexes commuting zones and t indexes month. The index m(t) is the month-of-year for month t.

We estimate one version of Equation 1 for each horizon $\ell \in \{1, \ldots, 6\}$. The regression tells

us how fluctuations in ENSO at time $t - \ell$ map into time t rainfall in each commuting zone. We refer to ℓ as the forecasting horizon, and we vary the forecasting horizon to determine the employment effects of rainfall forecasts with longer and shorter anticipation horizons. The employment effects are discussed and reported in the next section.

The key right-hand side variable is time $t - \ell$ ENSO interacted with CZ-specific coefficients. The coefficients capture regional variation in rainfall driven by ENSO. And these interaction terms are our excluded instruments in the second stage.

We identify the effect of ENSO on rainfall conditional on a number of controls. First, we include CZ by month-of-year fixed effects, $\rho_{c,m(t)}$, which subsume CZ fixed effects. These fixed effects condition out fixed features of the CZ as well as CZ-specific seasonal patterns that might otherwise lead to spurious relationships between rainfall and employment. We also include time fixed effects, which conditions out common time series patterns across the country. These fixed effects also condition out the time series pattern of ENSO, meaning that we are identified off of CZ-specific differences in the effect ENSO has on rainfall. The empirical specification can thus be viewed as a type of difference-in-differences estimator. The first stage identifies locations where rainfall is affected by ENSO at different forecast horizons, ℓ . For a given horizon, the locations where ENSO has a strong effect will be the "treatment group" while locations where ENSO does not have a strong effect are the "control group."

Two of the further controls are of particular importance. CZ-specific coefficients on lags of rainfall isolate variation driven by ENSO from other types of regional weather fluctuations. One could use past realizations of rainfall to forecast future rainfall, but this strategy could be invalid if the past rain realizations themselves have persistent effects on productivity. We control for rainfall up until time t to prevent contamination of our results by persistent effects of rainfall that occurs prior to time t. Second, CZ-specific coefficients interacted with the $\ell + 6$ month lag of ENSO isolate news about ENSO that has arrived in the last 6 months. This gives us relatively precise information about the timing of the arrival of information. Knowing when information is arriving helps us characterize the dynamics of adjustment.

Figure 1 shows the first stage estimates along with how they vary by forecast horizon. There is one point in the figure for each CZ in the sample. The points show the relationship between the CZ-specific ENSO coefficients from the version of Equation 1 estimated with the one-month lag of ENSO and the difference between the six-month lag coefficients and the one-month lag coefficients. In other words, the x-axis is the value of the γ_c coefficients from the $\ell = 1$ version of Equation 1. The y-axis is the difference between the $\ell = 6$ coefficients and the $\ell = 1$ coefficients.

The main take-away from the figure is that there is appreciable variation in the relationship between ENSO and rainfall across time for the different CZs. We exploit these differences to identify the effect of different anticipation horizons.



Figure 1: Differences in First-stage Coefficients

Notes: The x-axis of the figure shows coefficient estimates from the 1-month ahead $(\ell = 1)$ version of Equation 1. The y-axis is the difference between the 6-month ahead coefficients and the 1-month coefficients. There is one point for each CZ in our sample.

For example, focusing on rainfall responses at the 1-month horizon, there are hundreds of commuting zones where rainfall is close to non-responsive to ENSO (0 on the x-axis), and these are "control" CZ's. However, there are many more where ENSO increases rainfall, and even a handful where it *decreases* rainfall. At the extreme, in Pecos, TX, a one standard deviation increase in ENSO increases rainfall by 80 log points (120%).²⁰ This cross-sectional variation is one source of identification that we use.

However, we are also interested in how employment responses differ depending on the horizon of the forecast. This second source of identifying variation is shown by the y-axis. Consider, for example, Yuma, AZ, and Phoenix, AZ. These commuting zones are within the same state and are roughly equally responsive at the one-month horizon (.4, x-axis). But Yuma is twice as responsive at six months as it is at one month, while Phoenix is half as responsive. As a result of variation like this, two CZ's might experience the same short-run "ENSO treatment" but different long-run treatments, allowing us to separately identify

²⁰Figure A2 shows ENSO over time, measured in standard deviations. It is important to note that large and sudden shocks are not rare. For instance, in the late 1990's it rose by 2.5 standard deviations over less than a year, only to fall by 3.5 standard deviations a year later.

responses to short-run and long-run forecasts.

The figure also gives an indication of the locations where rainfall tends to be more strongly associated with ENSO. Stronger ENSO predicts rainfall increases consistently across the South and Southwest of the U.S. from Texas to California, up the West Coast, and to a lesser extent in areas around Georgia down to Florida. Stronger ENSO predicts lower rainfall in the northern Rocky Mountains and upper Midwest.

Figure 1 shows the variation in forecast coefficients and gives a measure of magnitude.²¹ One can get a further sense for the strength of the rainfall predictions of ENSO by examining the F-statistic from the first stage regression. Because of the large number of variables we estimate (one for each CZ), we calculate a conservative, lower bound on the conventional F-statistic using a version of Equation 1 where we interact ENSO with a state fixed effect rather than a CZ fixed effect (holding everything else about the data and estimation fixed). This is a lower bound because it throws away all within-state variation from Figure 1. The resulting Montiel Olea and Pflueger (2013) effective F-statistic is 50.3 at a 1-month horizon and 40.0 at a 6-month horizon, compared to a critical value for worst case bias of 5% of 26.7. At all forecast horizons, the ENSO-based predictions are strong and jointly significant.

4.2 Effect of predictable rainfall on employment

Our core interest is in employment adjustments in response to rainfall, depending on how far in advance that rainfall could be predicted. As discussed above, we rely on two sets of facts: first, that rainfall has important effects on construction sector productivity. Second, that ENSO fluctuations predict medium-term rainfall changes in the United States and that there is rich heterogeneity across place in how much and when these fluctuations translate into rainfall. This heterogeneity allows us to build an instrumental variables (IV) regression where we instrument for local rainfall using the CZ-specific ENSO response but are still able to control for CZ fixed effects (time-invariant geographic heterogeneity) and month fixed effects (arbitrary nationwide trends).

The second stage of the IV strategy is given by

$$\Delta_{\ell} \ln(Emp)_{c,t} = \delta_t^2 + \rho_{c,m(t)}^2 + \beta_{\ell} \widehat{\ln(Rain)}_{c,t}^{\ell} + \eta_c^2 ENSO_{t-\ell-6}$$

$$+ \theta_{c,1}^2 \ln(Rain)_{c,t-1} + \theta_{c,2}^2 \ln(Rain)_{c,t-2} + \sum_{k=3}^{12} \pi_k^2 \ln(Rain)_{c,t-k}$$

$$+ \sum_{k=1}^{12} \zeta_k^2 Temp_{c,t-k} + \varepsilon_{c,t}$$
(2)

²¹For the magnitude of the 6-month ahead forecast undifferenced, see Figure A4.

where $\Delta_{\ell} \ln(Emp)_{c,t}$ is $\ln(Emp)_{c,t} - \ln(Emp)_{c,t-\ell-1}$. Estimating the effect on the change in employment reduces unit root concerns and helps eliminate confounding variables up to time $t - \ell - 1$.²² Estimating with lags of employment on the right-hand side yields similar estimates but requires the addition of employment lags in the first stage regression, so we prefer the specification here.

The primary right-hand side variable of interest is $\widehat{\ln(Rain)}_{c,t}^{\ell}$. The variable is generated by the first stage regression, Equation 1, and is the expected rainfall in commuting zone cat time t which could be forecast from information available ℓ months beforehand.²³

The second stage includes every variable from the first stage except the excluded instruments: $ENSO_{t-\ell}$ interacted with CZ-specific indicators. In particular, it includes controls for lags of temperature (which improve precision), lags of realized rainfall, lags of ENSO interacted with CZ indicators, and fixed effects for the month and CZ by month-of-year. These controls remove confounding variation that is fixed within a location, that varies seasonally within a location, that varies over time nationwide (including the month-to-month variation in ENSO itself), and that might be caused by weather arriving before time t.

Table 2 displays estimated coefficients for the one through six-month ahead forecasts $(\ell \in \{1, ..., 6\})$ on employment at time t. There is a clear and monotonic pattern. Recall that rainfall substantially reduces construction productivity. Despite that, employment barely responds to rainfall increases that could only be predicted one month in advance. In contrast, employment is nearly five times as responsive to rainfall that could be predicted six months in advance, and the difference in response at the two different horizons is statistically significant at the 5% level.²⁴ The further ahead the increased rainfall could be predicted, the larger the employment response.

Why might employment respond more to forecasts available further in advance? One explanation is that labor market frictions and adjustment costs make instantaneous responses costly. Above, we discussed several of these costs, such as implicit layoff taxes generated by the financing rules of the unemployment insurance system.²⁵ We also discussed a series of costs in terms of recruiting, screening, and hiring. It is important to keep in mind that these hiring costs are relevant for firms even if our empirical results reflect employment *decreases*

 $^{^{22}}$ Formal panel unit root tests using the Im et al. (2003) procedure reject the null that the series contains a unit root, but the high degree of autocorrelation in employment still leads us to prefer a specification in differences.

²³The estimation is done through instrumental variables, so the hat on top of the rain variable is purely a notational reminder that the second stage regression involves unbiased rainfall *forecasts*.

²⁴While few of the other coefficients are statistically significantly different from one another, the monotonic pattern and difference in magnitudes is clear. The calibration described in Section 5 reflects the uncertainty in these parameter estimates.

 $^{^{25}}$ Again, even in the absence of rules like these, if sudden firing raises workers' perceptions of the riskiness of the job, then it can still be costly to firms who have to renegotiate higher wages to compensate workers for risk (as argued by Smith (1776)).

(1)	(2)	(3)	(4)	(5)	(6)
1	2	3	4	5	6
005* (.003)	010^{**} (.004)	015^{**} (.006)	017^{**} (.007)	019^{**} (.008)	026^{***} (.010)
.655 194,947	.707 194,947	.728 194,947	.730 194,947	.717 194,947	.689 194,947
	(1) 1 005* (.003) .655 194,947	$\begin{array}{ccc} (1) & (2) \\ 1 & 2 \\ \hline005^* &010^{**} \\ (.003) & (.004) \\ \hline .655 & .707 \\ 194,947 & 194,947 \end{array}$	$\begin{array}{c cccc} (1) & (2) & (3) \\ \hline 1 & 2 & 3 \\ \hline005^* &010^{**} &015^{**} \\ (.003) & (.004) & (.006) \\ \hline .655 & .707 & .728 \\ 194,947 & 194,947 & 194,947 \end{array}$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$

Table 2: Employment responses to predictable rainfall by length of forecast

* p < .10, ** p < .05, *** p < .01. Table displays estimated elasticities of the response of employment at time t to rainfall estimated at time $t - \ell$ to occur at time t. Standard errors clustered at the commuting zone (CZ) level are in parentheses. All estimates come from instrumental variables regressions that include time fixed effects, CZ fixed effects, and CZ-by-month-of-year fixed effects, as well as lagged rainfall, employment, and temperature controls shown in the estimating equations 1 (first stage) and 2 (second stage). Rainfall predictions are the primary explanatory variable and are based on CZspecific responses to changes in the El Niño/Southern Oscillation. See text for details.

in response to increased rain. This is because rainfall is only a temporary productivity shock. Firms will soon need to re-staff after excess rain subsides, and hiring new workers will require paying those costs.

Recognizing this, a firm which suddenly discovers that it will face a negative productivity shock during the next month (estimates for $\ell = 1$) might simply prefer to continue to pay the unproductive workers for a short time, rather than laying them off and having to replace them through a costly screening process soon after. After all, the costs of replacing a worker are roughly equal to one month salary (Dube et al., 2010), and it is unlikely that workers' productivity falls all the way to zero during the period of excess rain. A firm which discovers that it has six months until that productivity decline, on the other hand, can take advantage of the regular turnover process and simply delay hiring replacements for those who leave.

Whether or not this is realistic depends on the rate of natural turnover in construction, as well as the magnitude of the rain-induced employment responses that we document. The turnover rate can be measured using linked Current Population Survey (CPS) data, which shows that in a given month, 9% of construction workers leave their current employer for another. To understand the magnitudes of the responses that we identify, consider a very ENSO-responsive CZ that sees rainfall rise by 40 log points in response to a one standard deviation ENSO shock (the heterogeneity in ENSO-responsiveness is given in Figure 1). Even at the 6-month horizon, when employment shows the greatest responses, the second stage estimates in Table 2 suggest that this translates into a roughly 1% decline in employment.²⁶ It is not difficult to imagine a firm with one-month departure rates of 9% managing a 6-month decline of 1% exclusively using hiring delays. While one interpretation would suggest this means our empirical estimates are too small to be important, we disagree; for many firms, the opportunity to save or waste 1% of labor costs is a critical business decision, particularly in a competitive industry like construction.

These back-of-the-envelope calculations are useful for interpreting our reduced form estimates and the plausibility of the firm responses that we identify. By themselves, however, they do not allow us to understand the implications of rainfall shocks for worker and firm outcomes (aside from employment changes), and they do not allow us to derive quantitative implications of rising rainfall volatility driven by climate change. In the next section, we develop a model of employment dynamics that links them to adjustment costs for workers and firms. We use our reduced form estimates of the employment responses to calibrate the adjustment costs that firms face, and use this to quantify the consequences of rising rainfall volatility and the value of improving the quality of forecasts.

5 Model and counterfactuals

5.1 Model of labor market with adjustment costs

The elasticity estimates in Table 2 capture the response of the construction sector within a local labor market to news about rain. These responses reflect both the adjustment of the supply of labor in the construction industry as well as demand for labor by construction firms. In turn, adjustments on both sides of the market reflect conditions in other sectors of the economy. For example, construction services are intensively used for investment, so demand for construction services reflects investment decisions throughout the local labor market. To account for these cross-sector linkages, we model the entire local labor market's response to rain, incorporating spillover effects through the input-output structure of the economy. In the presentation of results, we focus on the construction industry for clarity, but the underlying model includes all sectors.

Accordingly, we model the local labor market as a small, open, multi-sector economy, populated by firms in each sector and a representative household. Firms use the output of all sectors as intermediate inputs, as well as to invest in capital formation. Households use the output of sectors for both consumption as well as to invest in housing. This setup allows us to directly calibrate the model to match make and use input-output tables from the Bureau of Economic Analysis (BEA).

 $^{26}\exp\{.4 \times -.026\} = .9896$

In our counterfactuals, we focus on the importance of different levels of adjustment costs for firms. We introduce these costs as employees diverted from production when a firm changes their hiring rate over time.

Specifically, the output of each firm in industry i at time t is Cobb-Douglas in capital, K_{it} ; labor used in production, L_{it} ; and an index for commodities used as materials, M_{it} .²⁷

$$Y_{it} = A_{it} \left(K_{it}^{\alpha_i} L_{it}^{1-\alpha_i} \right)^{1-\gamma_i} M_{it}^{\gamma_i}.$$
(3)

The variable A_{it} denotes total factor productivity. We assume that construction sector productivity depends negatively on realized rainfall with a constant elasticity, which we denote by ϵ .

Hiring is costly for the firm because it requires the use of employees devoted to the task. The productive labor input of the firm consists of those employees which it does not need to devote to hiring new workers. Given its current number of workers, N_{it} , the fraction of its labor force that must be devoted to hiring is

$$1 - \frac{L_{it}}{N_{it}} = \frac{\kappa}{2} \left(\frac{H_{it} - H_{i,t-1}}{H_{i,t-1}} \right)^2$$
(4)

If a firm in steady state $(H_{it} = H_{i,t-1})$ wants to increase or decrease its hiring rate, then it must divert some of its workforce away from production. These costs are convex so that larger deviations of H_{it} from $H_{i,t-1}$ incur greater costs, and κ is the key parameter that pins down the magnitude of these hiring costs.

The importance of hiring adjustment costs can be seen through the first order condition for new hires, which states that the marginal benefit of a new worker must equal the marginal cost of a new hire. The marginal benefit is forward looking because workers may stay with the firm for many periods. In particular, the firm's employment evolves as

$$N_{it} = (1-s)\Pi_{i,t}^{\text{Stay}} N_{i,t-1} + H_{it}$$
(5)

where $\Pi_{i,t}^{\text{Stay}}$ is the fraction of the firm's past employees who stay in the industry, and s is the fraction of these workers who separate from the firm and must find a new job in the industry. The first term captures the total number of workers who stay with the firm over time.

Given the law of motion in (5), the shadow value of another worker to the firm, V_{it} ,

²⁷The firm accumulates capital using a Cobb-Douglas technology in commodities, and, similarly, the materials index is Cobb-Douglas in commodities. See Appendix B.1 for details.

satisfies the following forward looking condition.

$$V_{it} = P_{it}^{Y} MPL_{it} \frac{L_{it}}{N_{it}} - W_{it} + \frac{1-s}{1+r} \mathbb{E}_t \Pi_{i,t+1}^{\text{Stay}} V_{i,t+1}$$
(6)

where P_{it}^{Y} denotes the competitive price of industry *i* output, $MPL_{it} \equiv (1 - \alpha_i)(1 - \gamma_i)\frac{Y_{it}}{L_{it}}$ is the marginal product of labor, W_{it} denotes the competitive wage in industry *i*, and *r* denotes the interest rate at which the firm is externally financed. The firm values a worker based on the expected discounted value of the gap between the marginal revenue product of labor and the wage, accounting for the chance that the worker stays with the firm over time. When the firm anticipates that future marginal revenue products will be high relative to wages, it will value building up its current labor force.

It will do so trading off these benefits against current hiring costs. That is, the value of an additional worker must equal the marginal cost of hiring an additional worker.

$$V_{it} = P_{it}^{Y} MPL_{it} \kappa \left(\frac{H_{it} - H_{i,t-1}}{H_{i,t-1}}\right) \frac{N_{it}}{H_{i,t-1}} - \mathbb{E}_t \frac{P_{i,t+1}^{Y} MPL_{i,t+1}}{1+r} \kappa \left(\frac{H_{i,t+1} - H_{it}}{H_{it}}\right) \frac{H_{i,t+1}}{H_{it}} \frac{N_{i,t+1}}{H_{it}}$$
(7)

The expression for the marginal cost has two terms. The first represents the cost to the firm from reduced output in the present due to diverting labor toward hiring, while the second represents the gain to the firm from avoiding future lost output. When the value of additional workers is high in the present, the firm will be more willing to forego present output to build up its labor force through hiring.

Notice that in the absence of hiring costs (when $\kappa = 0$), the firm will hire additional workers until the marginal benefit of hiring is zero. In this case, because $V_{it} = 0$, the valuation condition in (6) reduces to stating that the wage must equal the marginal revenue product of labor. Hiring adjustment costs lead to departures from the equalization of wages to the marginal revenue product of labor, and link hiring decisions to the anticipated future gaps between the marginal revenue product and wages. When the present value of these gaps is positive, the firm will tend to increase its hiring. News of reduced future productivity in the construction sector—such as increases in forecasts of rain—will reduce the value of hiring workers today, leading to current reductions in employment.

Of course, the total impact of reductions in future productivity depend also on the response of labor supply. Workers will also anticipate that they will face low future wages, and, if they also face adjustment costs, they will want to exit the industry preemptively to avoid being stuck with low income in the future.

To capture this possibility and maintain tractability, we close the model by introducing a representative household who provides all its members with perfect consumption and housing insurance, and allocates individuals to work across sectors. In addition to working, each individual can also be non-employed. Let i = 0 index this non-employment state, and denote individual n's employment status at time t as $i_t(n) \in \{0, 1, ..., M\}$ where M is the number of industries. Average utility of household members is

$$\mathcal{U}_{t} = \sum_{h=0}^{\infty} \beta^{h} \left[\ln C_{t+h} + \mu^{S} \ln S_{t+h} + \int_{0}^{1} \nu_{i_{t+h-1}(n), i_{t+h}(n), t+h}(n) \mathrm{d}n \right]$$
(8)

where $\beta \in [0, 1)$, C_t denotes a Cobb-Douglas consumption index, S_t denotes services from housing, μ^S captures the importance of housing services, and $\nu_{i,i',t}(n)$ is an individual preference shock for transitioning from i at t-1 to i' at time t. These preference shocks generate heterogeneity in how workers decide to switch sectors throughout their career. Changes in relative wages across sectors (driven by productivity shocks to the construction sector, for example) can drive these decisions, but these preference shocks allow workers to have disutility from leaving their job, or from switching sectors even in the absence of pecuniary benefits to doing so.

We assume that these preference shocks are Gumbel distributed independently across i' with shape $1/\theta$ and independently and identically distributed over time and individuals. The parameter θ captures heterogeneity across workers in their preferences to switch sectors. When θ is low, there is low dispersion across individuals so individuals are more willing to shift across sectors as the value of employment in each sector changes. In other words, when other factors like family needs, geographic mobility, injury, or aging or educational attainment (all of which we model in a reduced form way as a preference shock) are unimportant, then θ will be low implying little dispersion in preferences, and workers will be very responsive to changes in relative pay across sectors. We interpret this as suggesting very low adjustment costs (on the household side), because workers are very willing to respond to monetary incentives. If, on the other hand, θ is high, then dispersion in non-monetary preferences is large, and mobility responses to shocks to the relative wage will be muted.

For each moment in time, the representative household chooses consumption, housing investment, and an allocation of individuals across industries to maximize the expected discounted average utility of its members subject to its budget constraint and a housing law of motion. Here, we focus on the allocation of workers across industries. See Appendix B.1 for details on the household's consumption and investment decisions.

Under our Gumbel assumption on idiosyncratic preference shocks, the utility value of allocating another worker to employment status i during period t is

$$U_{it} = \frac{W_{it}/P_t^C}{C_t} + \beta \mathbb{E}_t \left[\theta \ln \sum_{i'=0}^M \omega_{i,i'} e^{U_{i',t+1}/\theta} \right]$$
(9)

where P_t^C is the price of consumption, W_{it} denotes income from employment status of iat time t—either the wage in a sector or income when non-employed—and $\omega_{i,i'} \geq 0$ with $\sum_{i'=0}^{M} \omega_{i,i'} = 1$ captures average preferences for transitioning from i to i'. The value of an additional worker in industry i depends on the marginal utility value of the real wage in that industry as well as anticipated future real wages across all industries, accounting for the chance of reallocating the worker in the future. In this way, labor supply is forward looking, with future wages impacting the current allocation of workers across industries.

Specifically, the share of workers shifting from i at time t - 1 to i' at time t is

$$\Pi_{i,i',t} = \frac{\omega_{i,i'} \exp(U_{i',t}/\theta)}{\sum_{\tilde{i}=0}^{M} \omega_{i,\tilde{i}} \exp(U_{\tilde{i},t}/\theta)}$$
(10)

leading to the following employment law of motion

$$N_{i't} = \sum_{i=0}^{M} \prod_{i,i',t} N_{i,t-1}.$$
(11)

This result provides an interpretation for $\omega_{i,i'}$ as the share of workers transitioning when all employment statuses have equal value (when $U_{i't} = U_{it}$ for all i', i). It also shows that $1/\theta$ is the semi-elasticity of employment flows to changes in the value of employment. When labor supply adjustment costs are high, it takes a larger change in the value of employment to generate a shift of workers across sectors.

Finally, local production of commodities arises from combining the output of each industry to match make input-output tables from BEA. We assume a Cobb-Douglas structure so that the price of each commodity is a Cobb-Douglas price index of local output prices. In turn, consumption, capital investment, and housing investment price indices are Cobb-Douglas in commodity prices with weights matching shares in the use input-output tables from the BEA. Further, we classify some commodities as non-tradeable (forcing local production to match local demand) based on observed imports and exports as a share of value added. See Appendix B.1 for details.

5.2 Calibration

We calibrate the model to match input-output tables from BEA, monthly employment transition rates across sectors from the Current Population Survey, and a 2% annual percentage rate for household discounting and the firm's interest rate. See Appendix B.2 for summary statistics of the calibration. This pins down all parameters except those controlling adjustment costs, the effect of rainfall on construction productivity, and the impulse response of rainfall to news about rain. For the response of rainfall to news about rain, we fit an AR(5) process to empirical estimates of the change in rainfall with respect to news about rainfall shocks (pooling across forecast horizons). In other words, we estimate the empirical impulse response of rainfall after time t to our forecasts of rainfall at time $t - \ell$ for $\ell = 1, \ldots, 6$. We then fit an AR(5) model to the set of these impulse response estimates and use the coefficients to account for the empirical persistence of rainfall effects. See Figure A7 in the appendix for the fitted response.

To pin down the scaling effect of the impact of rain on productivity, we choose the elasticity of construction productivity to rain, ϵ , to normalize the model's average prediction (across forecast horizon) for news about rain to match the average of our estimates in Figure 2. To estimate the labor supply and demand adjustment cost parameters, θ and κ , we minimize the distance between the model's normalized prediction for adjustment of construction employment to news about rain and our estimates in Table 2. Note that this effectively provides five target moments (six coefficients minus one absorbed degree of freedom) for the two labor adjustment cost parameters.

The model calibration and fit to the empirical employment elasticities allow us to estimate the two adjustment cost parameters in the model: θ and κ . Recall that θ measures the degree of heterogeneity in worker preferences for moving across sectors. A low value of θ would mean that workers are relatively homogeneous, and for a given change in the value of being in a sector, a large fraction of the workers would want to change their sector. High values imply that it would take a large shift in value to induce substantial movement of workers across sectors. Thus, we view θ as capturing labor supply adjustment costs.

Es	timates	Fit By Horizon							
	Value	h	0	-1	-2	-3	-4	-5	-6
θ	21.2	Uncentered \mathbb{R}^2	99.6%	98.2%	98.1%	96.4%	94.4%	94.5%	95.8%
κ	43.7	Centered \mathbb{R}^2	97.3%	93.8%	93.6%	89.3%	85.3%	83.9%	-
ϵ	0.11	l	1 to 6	1 to 6	2 to 6	3 to 6	4 to 6	$5~{\rm and}~6$	6
		Target Moments	\checkmark						

Table 3: Inferred Parameters and Model Fit

Notes: Parameter values minimize the mean squared error in the target moments between the estimates and the model with ϵ chosen to scale the average of model predictions to match the average of estimates across ℓ . We omit the centered R^2 at horizon 6 because there is only one comparison value so it is mechanically 100%.

The κ parameter measures demand side adjustment costs. It is the cost to the firm of adjusting its hiring rate, as shown in Equation 4. For a higher value of κ , firms must devote

more of their production toward hiring efforts (for instance by using more time for interviews, advertising positions, reviewing applications, or training).

The values for θ and κ that best fit the set of empirical rainfall forecast and employment elasticities are shown in Table 3. We can interpret the estimate of θ using Equation 9 and 10. An increase in the real wage difference between two sectors in a single month by 100% of steady state consumption (holding fixed expectations of the future value of both sectors) will increase the odds that a worker changes sectors by 5% during that month.²⁸ Using Equation 4, the estimate of κ implies that devoting 1% of a firm's labor force to hiring rather than production during a month will generate a 2.1% increase in the firm's hiring rate, while devoting 10% of the labor force generates a 6.8% increase in the hiring rate during that month.

The table also shows measures of model fit, both for the target moments—the elasticities of employment with respect to rainfall forecasts at the time when rainfall arrives—and for non-target moments of that same elasticity for months after news has arrived but before the rainfall shock has occurred. The empirical estimates for these moments can be found in Figure A6. The table reports R^2 values that are the amount of variation in the empirical point estimates that is explained by the model-derived employment elasticities. One can see that even for these non-targeted moments, the model fit is good, with a minimum R^2 of 84%.

5.3 Counterfactuals

One advantage of our model is that it allows us to interpret the rainfall shocks that we observe in more meaningful units. Using our benchmark calibration, we find that a one standard deviation above-mean rainfall shock is equivalent to an 11% loss of productivity. An 11% productivity decline is clearly meaningful, and in this section, we focus on how firms respond, how the costs of this shock are shared between workers and firms, and how the size of adjustment costs determines these responses.

Because our primary interest is in how firms might improve their ability to adapt (e.g., by negotiating more flexible labor contracts with workers or by diversifying their mix of projects to allow project-specific labor levels to be more responsive), we keep worker-side adjustment costs (θ) constant at the level of our baseline calibration. Instead, our counterfactuals focus on how changes in *firms*' ability to respond to shocks (κ) affects their employment dynamics and profits, and how this interacts with the quality of weather forecasts.

²⁸That is, $\ln \frac{\Pi_{i,i',t}}{\Pi_{i,i,t}}$ changes by .05 when $(W_{it} - W_{i't})/P_t^C$ increases by 100% of steady state consumption.

5.3.1 How firm-side adjustment costs drive employment dynamics

Above, we summarized a number of reasons why firms cannot freely or flexibly adjust their labor pool. An advantage of our method is that it captures these costs without imposing assumptions on the *types* of activities that are costly. Rather, it simply infers them from the dynamic pattern of employment responses. How would employment responses differ if the magnitude of adjustment costs was different? What features of the dynamic adjustments help us to infer those costs? How important are the magnitudes of responses, rather than their dynamic pattern, for quantifying these costs? Figure 2 helps answer these questions.

Specifically, in Figure 2, we compare the empirical, reduced form estimates (from Table 2) with those derived from the calibrated model. Comparing the filled circles (our reduced form estimates from above) with the hollow squares (derived from our baseline calibrations of κ and θ), our model very closely replicates the dynamic responses that we observe in the data.

Figure 2: Empirical and Model-Based Elasticity Estimates



Notes: The figure shows coefficient estimates of the effect of rainfall forecasts on employment from Table 2 compared to the model-derived estimates for the baseline adjustment cost parameters (see Table 3) and the estimates under counterfactual adjustment costs.

More interesting, however, Figure 2 also presents simulated dynamics of *alternative* calibrations in which we assume that the firm-side adjustment costs (κ) are ten times larger (hollow diamonds) or one-tenth as large (hollow triangles). Our primary result that responses are monotonically increasing in forecast horizon is generally true across all three calibrations. However, the magnitudes of these differences matter.

As one would expect, the higher the adjustment costs, the lower the overall level of adjustment that occurs. However the *ratio* of responses to 6-month forecasts compared to 1-month forecasts becomes much larger. With adjustment costs 10 times larger than those we estimate in the data, responses to 6-month-ahead forecasts are eight times larger than the 1-month responses. Compare this to a five-fold difference under our baseline adjustment cost estimates. When adjustment costs are much smaller than our baseline estimates, 6-month responses are less than three times larger than 1-month responses. Thus, the larger are the costs of adjusting, the less overall adjustment there is, but also the greater is the wedge between responses to short-run and medium-run forecasts. This is because as adjustment costs become larger, it becomes more important to spread them out over time. This effect arises from the convex adjustment costs that make large adjustments much more costly than small ones, and it is important for understanding how the burden of adjustment is split between workers and firms in our counterfactuals below.

5.3.2 How the burden is shared between workers and firms

To understand how workers and firms split the costs of productivity shocks, we calculate the firm's present value of profits and the worker's present value of income. We focus on how this present value (PV) changes at time $t - \ell$ when a new forecast of a rainfall shock occurring at time t becomes available. We focus on the PV at the time of the information $(t - \ell)$ to capture the full effect of the rainfall shock as well as any adjustments that occur before the rain actually arrives. Our reduced form estimates show substantial adjustment does occur prior to the arrival of rain (between $t - \ell$ and t).

Figure 3 shows how the PV of profits and labor earnings change when news of a rainfall shock becomes available. For both, we present different estimates depending on the firm-side adjustment costs (κ) and the forecast horizon (ℓ).

Panel (a) shows the profit effects. Focusing on our baseline calibration (solid blue circles), a rainfall shock forecast only one-month ahead leads to a substantial 0.6% decline in the PV of profits because firms will see a large productivity decline but have limited ability to adjust employment with such short notice. Remarkably, having this information six months in advance instead, they are able to offset almost the entire profit loss via adjustments during those six months.

The hollow red circles and hollow gray squares show analogous results for larger and smaller adjustment costs, respectively. When adjustment costs are very large, the profit losses are much more substantial, but interestingly, are less steeply related to the forecast horizon. This is because even at six months, firms ability to adapt to the productivity shock is relatively constrained, and it makes little difference whether the forecast is available six months or one month in advance. Even more surprising, when the adjustment costs are low, the firm's profits actually *increase* in response to the shock. This is because the firm can easily adjust its labor to offset the costs of the productivity decline, but given the structure of the multi-sector economy, the firm will benefit from a boost of pent-up demand after the shock has passed. This means that the fully flexible firm can avoid the costs of productivity shocks while enjoying all of the benefits.

Figure 3: Change in Present Value of Profit and Labor Value When News About Rain Arrives



Notes: The figure shows the present value of profit for firms (Panel 3a) and the present value of earnings for workers (Panel 3b) as a function of the forecast horizon and three different values of the firm-side adjustment cost parameter, κ : high (red with hollow circles), baseline (blue with filled circles), and low (gray with hollow squares) adjustment costs. All values are calculated at the time when news about the rain shock arrives (time $t - \ell$ in Equations (1) and (2)). For example, the 6-month-ahead forecast effect is the present value from 6 months prior to the arrival of rain.

Panel (a) showed that firms monotonically benefit from longer forecast horizons that afford them more time to adjust, and that more costly adjustment substantially reduces the present value of profits. Of course, the primary mechanism of adjustment is through management of their workforce and hiring. For this reason, Panel (b) shows that workers' experience the exact opposite patterns of firms. The longer the forecast horizon, the larger the decrease in the PV of earnings. This is because long anticipation horizons allow the firm to adjust more effectively and pass the costs of the rainfall shock onto workers.²⁹ In the baseline calibration that matches the data, the lost PV of a 6-month forecast are roughly twice the losses of a 1-month forecast.

 $^{^{29}}$ Interestingly, when adjustment costs are extremely high, workers are actually slightly *better off* due to a negative productivity shock because the firm is essentially unable to adjust to the shock, but the workers still benefit from the post-shock boost of pent-up demand.

This result illustrates that the adjustment cost parameters we estimate to match the data imply that firm-side adjustment costs are larger than the worker-side ones. To see this, note that any agent facing convex adjustment costs values earlier information because it allows her more time to smooth the adjustment and minimize the incurred costs. Indeed, if we dramatically increase the worker-side costs, we see that workers, too, begin to benefit from earlier information relative to later information (results available upon request). In our context, though, firms' mechanism of adjustment is through layoffs and reduced hiring, both of which shift the costs of the shock onto workers. Despite the fact that workers also face costs of switching sectors, they prefer *less advanced* information because avoiding the costs that firms shift onto them more than outweighs the inconvenience that they themselves face from short notice. In summary then, the fact that the estimates in Panel (b) are upward sloping (less negative at short horizons) is not an inherent feature of our model, but a result that illustrates that demand-side adjustment costs are the key drivers of our results.

5.3.3 Quantifying the costs and benefits of adjustment for the firm

In this setting, firms are adapting by adjusting their inputs, including labor, in response to news about upcoming shocks. Works adapt by changing sectors. Adaptation is costly due to rigidity in adjustment, but it also brings benefits in the form of reduced damage from rain shocks. The present value of profit and earnings at the time when news arrives shown in Figure 3 capture the net effect of both of these channels. In this section, we break down the different sides of adaptation, focusing on costs and benefits for firms.

The pure adjustment costs—the level of κ times the change in hiring in Equation 4 are straightforward. Changing hiring always comes with a cost, and the size of the change determines the scale of the cost. If a manager wants to adapt to an upcoming shock by making large changes in employment, they either need to pay a high cost for a quick change or spread their employment changes out over time so that they does not climb so high up their convex cost curve. The empirical estimates and model results in Figure 2 show that the magnitude of the employment change grows monotonically with increasing forecast horizon. When assessing the trade-off between higher costs and faster employment changes, the empirical results indicate that managers and firms are deciding that it is worth it to engage in more adjustment if they can spread the costs out over time.

Figure 4 shows impulse responses of the hiring wedge over time. The figure translates the employment changes we see empirically into costs—in units of productivity lost by devoting it to hiring—that the firm pays to adjust its hiring. The figure shows these costs each month leading up to, during, and after a rainfall event. The lines show costs for different forecast horizons, and the panels (a, b, and c) show the value for high, baseline, and low values of κ respectively.





Notes: The figure shows impulse responses of the share of the firm's labor force devoted to hiring as in Eq. 4 (in units of percent of their labor force). The panels show costs each month from 6 months prior to the arrival of rain until 6 months after. The lines are at different forecast horizons, with the lightest colored lines at the 6 month horizon and the darkest colored lines at the 1 months horizon. The three subpanels show how these costs differ by firm-side adjustment cost parameter, κ .

One can see that in the baseline (panel b), firms find it worthwhile to pay substantial adjustment costs at all forecast horizons, a result that is reflected in the employment changes in Figure 2. For larger adjustment costs, firms want to adjust their labor force, but high costs prevent them for making as substantial of changes, particularly at short forecast horizons. Within the horizons we investigate, labor adjustment and cost rise monotonically every month. In a world with high κ , relatively surprising rainfall events would see very little labor force adjustment and a low cost paid for adaptation. This is precisely because it is would not be worthwhile to engage in such costly behavior. In panel (c), one can see that there is also little cost paid for adjustment when κ is low, but now it is despite a large change in the labor force. Firms would get a high "bang for their buck" when adapting to a shock if adjustment was cheap.

Together, these figures tell an important story for climate adaptation and estimation of climate damages. If the world is one with high costs of adaptation, like the high κ case here, then people will not engage in much adaptation, and researchers will not find substantial evidence for adaptation in the empirical record. Firm would want to engage in adaptation in this case, if better information was available. The low κ case will be one with a high quantity of adaptation and a low cost of adaptation. In between these two extremes, *ex ante* behavior is important to take into account both in terms of the quantity and cost of adjustment.

The direct adjustment costs are only part of the adaptation cost story. Prior to the arrival of rainfall, firms are also affected by changes in labor supply as well as spillovers from sectoral linkages and trade. The complex interaction of these different factors can be summarized by comparing the present value of profit for the firm at the time when rainfall arrives (time t) versus at the time when news about rainfall arrives $(t - \ell)$. The difference in these values captures the effect on firms of the myriad *ex ante* adaptation actions engaged in both by construction firms and by all other market participants. Figure 5 Panel 5a shows this difference for the baseline calibration. The top line is the present value of profit when news arrives. The vertical difference between these two lines captures the total "cost" of actions taken before the rain shock realizes.



Figure 5: Benefit and Cost of Adapting Given 1-month Change in Forecast Horizon

(a) Calculating adaptation costs and benefits using profits

(b) Costs and benefits from 1-month forecast horizon changes

Notes: Panel 5a shows the present value of profit for firms as a function of the forecast horizon for the baseline calibration. The blue line with solid circles shows the present value when rain arrives (time t). The red line with hollow circles shows the present value when news arrives (time $t - \ell$). Both lines use the same discount factor. The vertical distance between the lines is the cost of all actions taken prior to the arrival of rain. The distance between two points on the time t line is the "marginal" benefit of a 1-month increase in forecast horizon. Panel 5b shows the marginal cost and marginal benefit values for each forecast horizon in our sample.

Comparing two points along the blue line, one can also estimate the benefit of improving the forecast horizon. As we discussed in the previous section, firms benefit more from information available further in advance by spreading the adjustment over a longer period, allowing them to engage in more adaptation at a lower cost. We refer to the difference between two points along the top curve as the "marginal" benefit of forecasts because it is the reduction in damage provided by a 1-month increase in the forecast horizon.

Finally, Panel 5b shows the marginal benefit of forecast horizon changes in the top curve

(red with hollow circles) and the marginal cost of adjustment in the bottom curve (blue with filled circles). The marginal cost is the difference in costs between two forecast horizons, as indicated on the x-axis. One can see that for all forecast horizons, the benefit of a 1-month increase in forecast horizon is greater than the cost. In fact, due to spillovers across sectors and changes in labor supply, marginal costs are actually negative for firms at the shortest forecast horizons. Marginal benefit and cost converge for longer horizons. If firms were able to choose their forecast horizon (and if forecasts were costless for them to produce), they would choose a longer horizon than the longest we consider here. Concretely, based on the net benefits illustrated in Figure 5b and given the 2% APR financing cost that we calibrate to, a firm would be willing to pay 3.6 months of typical profits to learn about rainfall 5 months sooner every month during the year.

5.4 Future climate projections

Unless greenhouse gas emissions are brought down substantially, future changes in the climate are projected to increase the volatility of rainfall around the world. The most recent, comprehensive climate projections under the high climate forcing scenario (CMIP6 SSP5-8.5) indicate that the standard deviation of rainfall will rise 12%, on average, in the continental U.S. between now and the end of the century (O'Neill et al., 2016). The rainfall distribution is bounded below by zero, so this increase in volatility will lead to a higher probability of heavier rainfall. Figure 6 shows rainfall volatility projections for every year between now and 2100.³⁰

Greater rainfall volatility will have two effects in our model. First, any shock—whether positive or negative—will require costly adjustment to either avoid damages or to take advantage of benefits. For firms, these costs are the hiring wedge in Equation 4 and shown in Figure 4. Second, there will be a higher probability of larger rainfall shocks conditional on one occurring, and as discussed above, this is particularly salient for rainfall which is bounded above 0.

The results in previous sections give the effects of a one standard deviation increase in rainfall on employment, productivity, profit, and earnings. If the standard deviation of rainfall increases from about 62 mm per month today to about 70 mm per month by 2100, how will that affect the construction industry and labor market? Simply extrapolating from our current estimates, holding everything else fixed, this will result in an additional loss in value added for the construction sector from a typical rainstorm of 0.02 percentage points. This magnitude might appear small, but we emphasize that is the loss, over-and-above losses already incurred in the current climate, from routine month-to-month rainfall. Given an total

 $^{^{30}}$ Figure A9 shows the projected standard deviation increase after residualizing on location and month-ofyear fixed effects. The growth in volatility is the same.





Notes: The figure shows the projected standard deviation of monthly rainfall (in mm) each year from the present until 2100 in CMIP6 SSP5-8.5. Each point is the month-to-month standard deviation of rainfall for each grid point in the CMIP6 projections averaged across the continental U.S. The raw standard deviation is debaised to match the sample average from our estimation sample.

value added for the construction industry in the U.S. of roughly \$1 trillion per year (BEA, 2019), an increase in losses of this magnitude could translate into tens of billions of dollars of damage to the economy.

Our estimates are that profit for construction firms will change by a similar magnitude, with firms losing around 0.01 percentage points more in present-value profit from a typical rain storm. By comparing profit losses in the baseline scenario (shown in Figure 3) and profit losses under projected climate change, we can estimate the change in forecast horizon necessary for a firm to offset this impact.

Profit losses get smaller as firms can plan further ahead. To gain enough flexibility to avoid the projected extra profit loss, firms would forecasts to arrive just over half-a-month further in advance. In other words, a firm that faced a 12% larger rainfall event would be left no worse off it it was also given a half-month-further-ahead forecast. So one way to think about the climate damages that we estimate is that firms would need to find a way to increase their planning horizon by 50%, if they are currently planning one month ahead, in order to offset the damage from a 1 standard deviation increase in rain.³¹

For longer baseline forecast horizons, the size of the increase falls. Figure 3 shows that

³¹In practice, our results indicate that firms are planning all the way out to a 6-month horizon, so this comparison is just illustrative.

at longer horizons, firms are able to largely avoid profit losses, so smaller forecast horizon increases are sufficient to avoid extra damages under projected climate change. Planning at such long horizons is inherent challenging, however, so a small increase in forecast horizon might still be hard to achieve. Figure A11 shows the forecast horizon improvement needed to offset projected profit losses for each of the forecast horizons considered in the paper.

Long-range weather forecasting has been an area of focus for climatologists and meteorologists, particularly since the early 1990s. Some success has been achieved in forecasts at a monthly or seasonal horizon. The use of ENSO signals to forecast seasonal rainfall and temperature, which we exploit in this paper, is one of the clearest success stories. But such forecasts are inherently difficult, and barring improvements to existing forecasts, more careful attention by construction project managers, or other adaptations to reduce rigidity in the labor market, the increased volatility projected by climate models will translated into economic costs.³²

Losses stemming from the interaction of labor adjustment costs and weather volatility constitute damages from climate change that are, as far as we are aware, novel in the climate economics literature. The damages do have an important antecedent in earlier work that emphasized the role of capital and agricultural input adjustment costs in overall climate damage (Quiggin and Horowitz, 2003, Kelly et al., 2005). The focus of that work, however, differs from ours in an important way: they focus on the costs paid along the transition path between equilibria while we focus on adjustment costs paid due routine weather shocks. Thus, one could distinguish between the "transition costs" pointed out by the previous literature and the "adjustment costs" that we investigate. Adjustment costs and attendant damages are important because they are paid routinely. Indeed, every time a weather shock realizes the economy either pays adjustment costs or suffers damage from the change (even if counterfactually from not adjusting to capture a gain from beneficial shocks).

6 Conclusion

Climate change is expected to cause substantial damage to the global economy. Our understanding of that damage comes primarily from integrated assessment models (IAMs) that estimate the equilibrium response to a change in the climate or from microeconometric estimates of the acute effects of weather. Both of these sources overlook important dynamics. In this paper, we studied two understudied aspects of the economics of climate change where dynamics play an important role: labor market adjustment and the effects of rainfall volatility. We focused on these issues in construction, an economically important, climate

³²Toth and Buizza (2019) summarize the history and possible future efforts to achieve gains in monthly, seasonal, or even longer-horizon forecasts. More pessimistically, Scher and Messori (2019) have recently argued that climate change might make rainfall forecasting even harder in the future.

exposed industry that has itself been largely overlooked by previous research.

We found evidence that construction labor markets respond sluggishly to forecasts of rainfall shocks—an indication that market participants face adjustment costs. Empirically, unexpected rainfall is associated with little change in construction employment. Rainfall that can be anticipated well in advance, in contrast, leads to large changes in employment. Calibrating a multi-sector model labor market model we found that the labor market responses to forecastable rainfall imply large labor adjustment costs.

Adjustment costs are a source of climate damage, a reason for *ex ante* adaptation by market participants, and a limit to the adaptation process. Given the adjustment costs we estimate, weather shocks cause not only acute damage, but also involve substantial payment for adaptation. Existing methods that estimate the effects of weather on economic outcomes will miss this source of damage. Increases in weather volatility—which is already apparent in the recent historical rainfall record and is projected to continue under unmitigated climate change—will also imply higher climate damage due to more frequent or larger need for adjustment. This source of damage is overlooked by IAMs. In our context, both of these effects strengthen the case for public policy to reduce the emissions of climate pollutants.

For an individual firm facing a world of increasing climate volatility and costly adjustment, we showed that if firms can implement improvements to their hiring that bring down the cost of employment adjustment, then they can substantially offset the negative effects of weather variation. Firms can also take action to better incorporate longer-horizon information into their decision-making. For construction firms that must make multi-month contracts based on weather expectations, investing in innovations and process improvements that will allow this type of planning will be of high value to improve the flexible and resilience in the face of a changing climate.

References

- Abraham, K. G. and S. K. Taylor (1996). Firms' use of outside contractors: Theory and evidence. *Journal of labor economics* 14(3), 394–424.
- Ackerberg, D. A., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica* 83(6), 2411–2451.
- Adhvaryu, A., A. V. Chari, and S. Sharma (2013). Firing costs and flexibility: evidence from firms' employment responses to shocks in india. *Review of Economics and Statistics* 95(3), 725–740.
- Asker, J., A. Collard-Wexler, and J. De Loecker (2014). Dynamic inputs and resource (mis) allocation. Journal of Political Economy 122(5), 1013–1063.
- Ballesteros-Pérez, P., S. T. Smith, J. G. Lloyd-Papworth, and P. Cooke (2018). Incorporating the effect of weather in construction scheduling and management with sine wave curves: application in the united kingdom. *Construction Management and Economics* 36(12), 666–682.
- BEA (2019). Industry economic account data: GDP by industry. Accessed: 2020-11-01.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica* 77(3), 623–685.
- Burks, S. V., B. Cowgill, M. Hoffman, and M. Housman (2015). The value of hiring through employee referrals. *The Quarterly Journal of Economics* 130(2), 805–839.
- CIESIN (2017). U.S. Census grids 2010. https://sedac.ciesin.columbia.edu/data/ collection/gpw-v4. Accessed: 2018-08-30.
- Collard-Wexler, A. (2013). Demand fluctuations in the ready-mix concrete industry. *Econo*metrica 81(3), 1003–1037.
- Colmer, J. (2020). Temperature, labor reallocation and industrial production: evidence from india. pp. 1–28.
- Dube, A., E. Freeman, and M. Reich (2010). Employee replacement costs. IRLE Working Paper 201-10.
- ECMWF (2019). Long range. https://www.ecmwf.int/en/forecasts/ documentation-and-support/long-range. Accessed: 2019-10-27.
- Edin, P.-A., P. Fredriksson, M. Nybom, and B. Ockert (2017). The rising return to noncognitive skill. *IZA Discussion Paper*.

- Environmental Protection Agency (2009). Construction and development effluent guidelines and standards.
- Fankhaeser, S., F. Schlleier, and N. Stern (2008). Climate change, innovation and jobs. Climate policy 8(4), 421–429.
- Fishman, R. (2016). More uneven distributions overturn benefits of higher precipitation for crop yields. *Environmental Research Letters* 11(2), 024004.
- Fried, S. (2019). Seawalls and stilts: A quantitative macro study of climate adaptation. Society for Economic Dynamics, 44.
- Ginja, R., A. Karimi, and P. Xiao (2020). Employer responses to family leave programs. Working Paper.
- Graff Zivin, J. and M. Neidell (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics* 32(1), 1–26.
- Halpert, M. S. and C. F. Ropelewski (1992). Surface temperature patterns associated with the southern oscillation. *Journal of Climate* 5(6), 577–593.
- Halsey, E. (2016). How el nio will affect the construction industry. *Construction Business Owner*. Accessed: 2020-03-17.
- Im, K. S., M. H. Pesaran, and Y. Shin (2003). Testing for unit roots in heterogeneous panels. Journal of econometrics 115(1), 53–74.
- Isen, A. (2013). Dying to know: Are workers paid their marginal product? Working Paper.
- Jäger, S. and J. Heining (2019). How substitutable are workers? evidence from worker deaths. *Working Paper*.
- Jain, A., R. O'Sullivan, and V. Taraz (2020). Temperature and economic activity: Evidence from india. *Journal of Environmental Economics and Policy* 9(4), 430–446.
- Kahn, L. B. and F. Lange (2014). Employer learning, productivity, and the earnings distribution: Evidence from performance measures. The Review of Economic Studies 81(4), 1575–1613.
- Kellogg, R. (2014). The effect of uncertainty on investment: Evidence from texas oil drilling. American Economic Review 104(6), 1698–1734.
- Kelly, D. L., C. D. Kolstad, and G. T. Mitchell (2005). Adjustment costs from environmental change. *Journal of Environmental Economics and Management* 50(3), 468–495.

- Krane, S. and W. Wascher (1999). The cyclical sensitivity of seasonality in us employment. Journal of Monetary Economics 44(3), 523–553.
- Lemoine, D. (2018). Estimating the consequences of climate change from variation in weather. *NBER Working Paper* (No. 25008).
- Montiel Olea, J. L. and C. Pflueger (2013). A robust test for weak instruments. *Journal of Business and Economic Statistics* 31(3), 358–369.
- Naidu, S. and N. Yuchtman (2018). Labor market institutions in the gilded age of american economic history. In L. Cain, P. Fishback, and P. Rhode (Eds.), *The Oxford Handbook of American Economic History, vol. 1*, Chapter 14. Oxford: Oxford University Press.
- Oi, W. Y. (1962). Labor as a quasi-fixed factor. Journal of political economy 70(6), 538–555.
- O'Neill, B. C., C. Tebaldi, D. P. van Vuuren, V. Eyring, P. Friedlingstein, G. Hurtt, R. Knutti, E. Kriegler, J.-F. Lamarque, J. Lowe, G. A. Meehl, R. Moss, K. Riahi, and B. M. Sanderson (2016). The scenario model intercomparison project (scenariomip) for cmip6. *Geoscientific Model Development* 9(9), 3461–3482.
- Organization, I. L. (2016). Non-standard employment around the world: Understanding challenges, shaping prospects.
- Pendergrass, A. G., R. Knutti, F. Lehner, C. Deser, and B. M. Sanderson (2017). Precipitation variability increases in a warmer climate. *Scientific Reports* 7(1), 17966.
- PRISM Climate Group (2004). http://prism.oregonstate.edu.
- Quiggin, J. and J. Horowitz (2003). Costs of adjustment to climate change. Australian Journal of Agricultural and Resource Economics 47(4), 429–446.
- Ratner, D. (2013). Unemployment insurance experience rating and labor market dynamics. *FEDS Working Paper*.
- Ropelewski, C. F. and M. S. Halpert (1987). Global and regional scale precipitation patterns associated with the el niño/southern oscillation. *Monthly Weather Review* 115(8), 1606–1626.
- Rosenzweig, M. R. and C. Udry (2014). Rainfall forecasts, weather, and wages over the agricultural production cycle. *American Economic Review Papers & Proceedings* 104(5), 278–283.
- Scher, S. and G. Messori (2019). How global warming changes the difficulty of synoptic weather forecasting. *Geophysical Research Letters* 46(5), 2931–2939.

- Serfling, M. (2016). Firing costs and capital structure decisions. The Journal of Finance 71(5), 2239–2286.
- Shrader, J. (2020). Expectations and adaptation to environmental risks. *Working Paper*, 75.
- Smith, A. (1776). The Wealth of Nations: An inquiry into the nature and causes of the Wealth of Nations. Harriman House Limited.
- Toth, Z. and R. Buizza (2019). Weather forecasting: What sets the forecast skill horizon? In *Sub-Seasonal to Seasonal Prediction*, pp. 17–45. Elsevier.
- Trauner, T. J., S. Lowe, M. F. Nagata, and W. A. Manginelli (2018). Construction Delays (3rd ed.). Elsevier.
- UNECE (2019). Share of construction in GDP. Accessed: 2020-11-01.
- Williams, J. (2015). How the super El Nino of 1982-83 kept itself a secret. https://www.washingtonpost.com/news/capital-weather-gang/wp/2015/06/ 12/how-the-super-el-nino-of-1982-83-kept-itself-a-secret/. Accessed: 2019-11-01.
- Zhang, P., J. Zhang, and M. Chen (2017). Economic impacts of climate change on agriculture: The importance of additional climatic variables other than temperature and precipitation. *Journal of Environmental Economics and Management* 83, 8–31.

Appendix for online publication

A Supplementary figures and tables



Figure A1: Commuting Zones Where Construction Employment is Disclosed

Notes: The map shows commuting zones in the continental U.S. where construction employment is disclosed (white) versus suppressed (black) for all months in our sample.

Figure A2: ENSO Temperature Anomalies



Notes: The figure shows monthly average temperature anomalies in the equatorial Pacific Ocean as measured by the Niño 3.4 index. Gray bars are NBER recession dates.

Figure A3: CZ-average rainfall and temperature over 1990 to 2016 period



(a) Pop. weighted precipitation

(b) Pop. weighted avg. temperature

Panel (a) shows the time series average across the full sample (1990 to 2016) of population weighted monthly total precipitation (in mm) in each commuting zone. Panel (b) shows the population-weighted average temperature (in $^{\circ}$ C) for each commuting zone over the same period.

Figure A4: Comparison of First-stage Coefficients: $\ell = 1$ and $\ell = 6$



Notes: The figure shows coefficient estimates from Equation 1 for 1 and 6-month ahead predictions of the rainfall using El Niño/Southern Oscillation. Example locations where the predictions at the two horizons differ substantially are labelled. The dashed line is a 45° line.



Figure A5: Main Results Using Rainy Days

Notes: The figure shows coefficient estimates from Equation 2 for rainfall predicted 1 to 6 months ahead. Predictions are based on CZ-specific responses to changes in the El Niño/Southern Oscillation. In our primary specification (see Table 2), the core independent variable is the log number of millimeters of precipitation in the month. In this specification, the core independent variable is instead the number of days in the month with positive precipitation.



Figure A6: Effect of Rainfall Forecasts on Employment Prior to the Arrival of the Rain Shock

Notes: The figure shows coefficient estimates from Equation 2 for rainfall predicted 1 to 6 months ahead. Predictions are based on CZ-specific responses to changes in the El Niño/Southern Oscillation. The dependent variable is log employment measured h months before the rain was forecast to arrive. The bars are 95% confidence intervals.

B Model and Calibration

B.1 Model of labor market with adjustment costs

Average utility of household members:

$$\mathcal{U}_{t} = \sum_{h=0}^{\infty} \beta^{h} \left[\ln C_{t+h} + \mu^{S} \ln S_{t+h} + \int_{0}^{1} \nu_{i_{t+h-1}(n), i_{t+h}(n), t+h}(n) \mathrm{d}n \right] \quad \text{where} \quad C_{t} \equiv \prod_{j=1}^{J} \left(\frac{C_{jt}}{\mu_{j}^{C}} \right)^{\mu_{j}^{C}}$$
(12)

and $\nu_{i,i',t+1}(n)$ is Gumbel distributed independently across i' with shape $1/\theta$ and independently and identically distributed over time and individuals.

The representative household does not have access to external financing, but can use investment in housing to smooth consumption over time. We view this assumption as capturing how housing is a primary savings instrument for households in the U.S. economy.³³ It allows us to capture general equilibrium effects within a local labor market that impact residential investment decisions, while maintaining the convenience of a representative household.³⁴

Housing investment technology:

$$S_t = (1 - \delta^S) S_{t-1} + I_t^S \quad \text{where} \quad I_t^S \equiv \prod_{j=1}^J \left(\frac{I_{jt}^S}{\mu_j^S}\right)^{\mu_j^S}.$$
 (13)

Household budget constraint

$$\sum_{j=1}^{J} P_{jt}(C_{jt} + I_{jt}^S) \le \sum_{i=0}^{M} W_{it}N_{it} + T_t$$
(14)

where T_t denotes net taxes and transfers from the government.

The solution to the consumption and investment portion of the representative household's problem can be characterized as follows. The costs to the household of consumption and housing at time t are, respectively,

$$P_t^C = \prod_{j=1}^J P_{jt}^{\mu_j^C}$$
 and $P_t^S = \prod_{j=1}^J P_{jt}^{\mu_j^S}$ (15)

³³For instance, we get this assumption as a result if the representative household has access to external financing at an interest rate less than $1/\beta - 1$ but faces a no-borrowing constraint. In this case, they would like to borrow an infinite amount in steady state.

³⁴Note also that this setup lets us avoid the indeterminacy of steady state savings in a small open economy with a representative agent that can save at an exogenous interest rate equal to $1/\beta - 1$.

and consumption and investment in commodities satisfy, respectively,

$$P_{jt}C_{jt} = \mu_j^C P_t^C C_t$$
 and $P_{jt}I_{jt}^S = \mu_j^S P_t^S I_t^S$. (16)

That is, the household has constant expenditure shares across commodities for the purposes of consumption and residential investment. This result allows us to calibrate household consumption preferences and the housing investment technology using data on each commodity's share of total consumption and residential investment expenditure.

Investment in housing is characterized in terms of these price indices by

$$\mu^{S} \frac{C_{t}}{S_{t}} = \frac{P_{t}^{S}}{P_{t}^{C}} - \beta (1 - \delta^{S}) \mathbb{E}_{t} \frac{C_{t}}{C_{t+1}} \frac{P_{t+1}^{S}}{P_{t+1}^{C}}.$$
(17)

This condition states that the marginal rate of substitution between consumption and housing services has to equal the user cost of an additional unit of housing in units of consumption. The later can be interpreted as the implicit rental price of the housing stock in units of consumption. As a consequence, we can calibrate μ^S to the ratio of implicit rent on housing services to total consumption expenditure.

From our Gumbel assumption, the solution to the allocation of workers is as follows. The utility value of allocating another worker to employment status i:

$$U_{it} = \frac{W_{it}/P_t^C}{C_t} + \beta \mathbb{E}_t \left[\theta \ln \sum_{i'=0}^M \omega_{i,i'} e^{U_{i',t+1}/\theta} \right]$$
(18)

Employment law of motion:

$$N_{i't} = \sum_{i=0}^{M} \prod_{i,i',t} N_{i,t-1}.$$
(19)

where the share of workers shifting from i at time t - 1 to i' at time t is

$$\Pi_{i,i',t} = \frac{\omega_{i,i'} \exp(U_{i',t}/\theta)}{\sum_{\tilde{i}=0}^{M} \omega_{i,\tilde{i}} \exp(U_{\tilde{i},t}/\theta)}.$$
(20)

Turning to firms, the labor force in sector i, N_{it} , evolves as

$$N_{it} = (1-s)\Pi_{i,t}^{\text{Stay}} N_{i,t-1} + H_{it}$$

where $\Pi_{i,t}^{\text{Stay}} \equiv \Pi_{i,i,t}$.

The production function is

$$Y_{it} = A_{it} \left(K_{it}^{\alpha_i} L_{it}^{1-\alpha_i} \right)^{1-\gamma_i} \prod_{j=1}^J \left(\frac{X_{ijt}}{\Gamma_{ij}} \right)^{\Gamma_{ij}}$$
(21)

where $\gamma_i \equiv \sum_{j=1}^{J} \Gamma_{ij}$ and

$$L_{it} = \left[1 - \frac{\kappa}{2} \left(\frac{H_{it}}{H_{i,t-1}} - 1\right)^2\right] N_{it}.$$

Capital investment technology:

$$K_{i,t+1} = (1 - \delta^K) K_{it} + I_{it}^K, \text{ and } I_{it}^K \equiv \prod_{j=1}^J (I_{ijt}^K)^{\mu_j^K}.$$
 (22)

Here, δ^{K} denotes the depreciation rate of capital.

The firm's intermediate input expenditure and value added:

$$P_{jt}X_{ijt} = \Gamma_{ij}P_{it}^Y Y_{it}$$
 and $VA_{it} = (1 - \gamma_s)P_{it}^Y Y_{it}$

Letting R_{it} denote profit per unit of capital, the firm's profit and the wage bill are:

$$R_{it}K_{it} = \alpha_i V A_{it}$$
 and $W_{it}N_{it} = (1 - \alpha_i)V A_{it}$.

The firm's capital investment price index and investment expenditure:

$$P_t^K = \prod_{j=1}^J P_{jt}^{\mu_j^K}$$
 and $P_{jt}I_{ijt}^K = \mu_j^M P_t^K I_{it}^K$.

Firm investment no-arbitrage condition:

$$\frac{1}{1+r}\mathbb{E}_t R_{i,t+1} = P_t^K - \frac{1-\delta^K}{1+r}\mathbb{E}_t P_{t+1}^K$$

The shadow value of another worker to the firm is

$$V_{it} = P_{it}^{Y} M P L_{it} \frac{L_{it}}{N_{it}} - W_{it} + \frac{1-s}{1+r} \mathbb{E}_t \Pi_{i,t+1}^{\text{Stay}} V_{i,t+1}$$
(23)

where

$$MPL_{it} \equiv (1 - \alpha_i)(1 - \gamma_i)\frac{Y_{it}}{L_{it}}$$

Optimal hiring:

$$V_{it} = P_{it}^{Y} MPL_{it} \kappa \left(\frac{H_{it} - H_{i,t-1}}{H_{i,t-1}}\right) \frac{N_{it}}{H_{i,t-1}} - \mathbb{E}_t \frac{P_{i,t+1}^{Y} MPL_{i,t+1}}{1+r} \kappa \left(\frac{H_{i,t+1} - H_{it}}{H_{it}}\right) \frac{H_{i,t+1}}{H_{it}} \frac{N_{i,t+1}}{H_{it}} \frac{N_{i,$$

Local sector output gets combined according to a Cobb-Douglas technology in order to produce local commodity output. This links the local prices of sector output and commodities as

$$P_{jt} = \prod_{i=1}^{M} (P_{it}^Y)^{\Phi_{ij}}$$

Expenditure on local production in sector i equal to

$$P_{it}^{Y}Y_{it} = \sum_{j=1}^{J} \Phi_{ij} \left[P_{jt} \left(\sum_{i=1}^{M} X_{ijt} + C_{jt} + I_{jt}^{S} + \sum_{i=1}^{M} I_{ijt}^{K} \right) + G_{j} + NX_{jt} \right].$$

The market clearing condition for the local labor market is

 $P_{jt} = Z_{jt}$ if if j is traded and $NX_{jt} = 0$ if j is not traded,

where Z_{jt} is the world price for a traded commodity.

B.2 Calibration

Table A1: Calibration of Non-Sector-Specific Parameters

Parameter	Value	Source/Target
r	$1.02^{1/12} - 1$	2% APR
β	$1.02^{-1/12}$	$2\%~{ m APR}$
μ^S	.171	Housing to consumption expenditure (BEA use IO table)
δ^S	0.00154	Housing investment to housing expenditure (BEA use IO table)
δ^K	0.00145	Capital investment to payments to capital (BEA use IO table)
s	0.066	Share of workers in same sector with a new employer (CPS)

Notes: Payments to capital are inferred as value added net of labor compensation.

We use input-output tables from BEA and employment transition probabilities from the CPS to calibrate the model at the 2-digit NAICS level. Rather than reporting the full tables, we report aggregated values below to highlight the overall structure of the economy.



Figure A7: Fit of AR(5) Rain Process

Notes: The figure shows coefficient estimates and point-wise standard error bands for estimates of the impact of news about rain on realized rain by forecast horizon (blue lines with grey bands), as well as the fit of an AR(5) process to these estimates (dashed black line). We minimize the distance between the estimates (pooling across forecast horizons) and the prediction from an AR(5) process, using standard errors as weights.

Commodity/	μ_j^C	μ_j^S	μ_j^K	G_j	Capital	Comp./worker
Sector	, T				share (α_i)	share (A_s)
Construction	0.0039	0.7709	0.1788	0.1066	0.1887	6,441,050
Real Estate	0.0207	0.1255	0.0032	0.0	0.3425	$5,\!697,\!540$
Services	0.487	0.0143	0.0284	0.0	0.2008	$4,\!845,\!990$
Traded	0.4418	0.0893	0.7753	0.1168	0.1893	10,013,900
Traded (Residual)	0.035	0.0	0.0142	0.0	0.3265	4,057,720
Government/Misc	0.0117	0.0	0.0	0.7766	0.1415	$5,\!216,\!890$
Non-Employment						4,777,750

Table A2: Commodity/Sector Specific Calibrated Parameters/Targets

Notes: Data is from BEA use IO table. Compensation per worker share is in millions of dollars. Non-employment compensation is inferred as the benefit necessary to cover consumption expenditure and housing investment given labor compensation in other sectors.

Commodity/Sector	Const.	R.E.	Serv.	Traded	T.R	N.E.
Construction	0.904	0.0014	0.011	0.0201	0.0038	0.0597
Real Estate	0.0052	0.9148	0.0163	0.0178	0.0011	0.0448
Services	0.0017	0.0006	0.9329	0.014	0.0012	0.0496
Traded	0.0035	0.0008	0.0159	0.945	0.0018	0.0331
Traded (Residual)	0.0064	0.0006	0.0125	0.0169	0.8982	0.0653
Non-Employment	0.0252	0.0046	0.1483	0.0777	0.0166	0.7277

Table A3: Calibration target for $\omega_{ii'}$, sector-to-sector labor flows

Notes: The entries of each row are the share of workers moving from that sector to each other sector per month in the current population survey (CPS).

Table A4: Calibration of Γ_{ij} from BEA use IO table

Commodity/Sector	Const.	R.E.	Serv.	Traded	T.R	Govt./Misc.
Construction	0.0002	0.0338	0.0018	0.0028	0.0073	0.0181
Real Estate	0.0237	0.0929	0.0523	0.0272	0.0474	0.0179
Services	0.0702	0.1605	0.0674	0.0439	0.0275	0.0425
Traded	0.3695	0.2111	0.2207	0.3925	0.2525	0.2201
Traded (Residual)	0.0148	0.057	0.0174	0.0467	0.1916	0.0149
Government/Misc.	0.0046	0.0056	0.0078	0.0116	0.0078	0.009

Notes: Each column shows shares of that sector's revenue spent across commodities.

Table A5: Calibration of Φ_{ij} from BEA make IO table

Commodity/Sector	Const.	R.E.	Serv.	Traded	T.R	Govt./Misc.
Construction	0.9588	0.0012	0.009	0.0126	0.0076	0.0109
Real Estate	0.0	0.872	0.0178	0.0978	0.0	0.0125
Services	0.0	0.0011	0.931	0.0138	0.0003	0.0538
Traded	0.0002	0.0001	0.0054	0.9822	0.0024	0.0098
Traded (Residual)	0.0	0.0	0.0004	0.0036	0.9152	0.0808
Government/Misc.	0.0	0.0	0.0007	0.0019	0.0005	0.9969

Notes: Each row shows shares of that commodity's expenditure across sectors.



Figure A8: Present Value of Profit and Labor Value When Rain Arrives

Notes: The figure shows the present value of profit for firms (Panel A8a) and present value of earnings for workers (Panel A8b) as a function of the forecast horizon and three different values of the firm-side adjustment cost parameter, κ : high (red with hollow circles), baseline (blue with filled circles), and low (gray with hollow squares) adjustment costs. All values are calculated at the time when the rain shock arrives (time t in Equation (2)).

Figure A9: Projected monthly rainfall standard deviation for the continental U.S.



Panel (a) shows the projected standard deviation of monthly rainfall (in mm) each year from the present until 2100 in CMIP6 SSP5-8.5 (O'Neill et al., 2016). Each point is calculated by taking the month-to-month standard deviation of rainfall for each grid point in the CMIP6 projections then averaging those values across the continental U.S. The raw standard deviation is debaised to match the sample average from our estimation sample (by adding 2.04 to the projection values). Panel (b) shows the same standard deviations but where the monthly rainfall in each grid point is first residualized on month, year, grid point, and climate model fixed effects then the standard deviation is calculated.

Figure A10: Projected change in monthly rainfall standard deviation for the continental U.S.



The figure shows the projected growth in the standard deviation of monthly rainfall (in mm) each year from the present until 2100 in CMIP6 SSP5-8.5 (O'Neill et al., 2016). Each line is indexed to 1 in 2015. The gray lines are from different climate models in the CMIP database. The black line is the monthly average across models.

Figure A11: Forecast horizon improvement needed to offset projected losses



The figure shows the gain in forecasts needed to offset losses from the projected change in the standard deviation of monthly rainfall caused by climate change (shown in Figure 6) in our baseline calibration. The line is the horizontal distance between the baseline profit loss and the profit loss under the projected increase in rainfall volatility as a function of forecast horizon.