Improving Climate Damage Estimates by Accounting for Adaptation

Jeffrey G. Shrader*

June 5, 2023

Abstract

Climate change is projected to severely damage the global economy. Adaptation in response to the changing climate will affect how much damage ultimately occurs. An important source of uncertainty in existing damage estimates is the extent to which they include or exclude such adaptation. This paper shows how to identify damages by estimating economic responses to climate shocks while controlling for weather forecasts. The resulting empirical strategy also provides estimates of the benefits from forward-looking adaptation. The strategy is applied to study damages from climate shocks and adaptation benefits using detailed, firm-level data on commercial fishing and a novel dataset of climate forecasts. Without accounting for adaptation, direct damage estimates are substantially biased. Adaptation also yields large benefits, with forecasts allowing firms to time entry into the fishery to best avoid adverse conditions. (JEL:D22,D84,Q22,Q54)

*Columbia University, 420 West 118th Street, New York NY 10027 (e-mail: jgs2103@columbia.edu). Previous versions of this paper were circulated with the title "Expectations and Adaptation to Environmental Risks." Special thanks go to Debra Losey from the NOAA Fisheries library for helping me locate historical ENSO forecasts, to Anthony Barnston for providing me with the database of IRI ENSO forecasts, and to Dale Squires for providing access to fisheries data and for many helpful comments on drafts of the paper. Additionally, I would like to thank Roy Allen, Richard Carson, Chris Costello, Julie Cullen, Gordon Dahl, Olivier Deschenes, Matthew Gibson, Joshua Graff Zivin, Solomon Hsiang, Mark Jacobsen, Derek Lemoine, Kyle Meng, Ben Miller, Frances Moore, David Pierce, Mark Rosenzweig, Wolfram Schlenker, and Hendrik Wolff as well as seminar participants at the AAEA annual meeting, the ASSA Annual Meeting, Berkeley ARE, Columbia University, NBER EEE, the Occasional Workshop in Environmental Economics, UC San Diego, UC Santa Barbara, the Sea Grant Annual Meeting, and Scripps Institute of Oceanography for providing helpful feedback. This publication is supported by funding from NOAA Grant No. NA14OAR4170163, California Sea Grant College Program Project No. E/MRE-8, and the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1144086. Any opinion, findings, conclusions, or recommendations expressed in this material are those of the author and do not necessarily reflect the views of California Sea Grant, state agencies, NOAA, the U.S. Dept. of Commerce, or the National Science Foundation.

1 Introduction

How much damage will climate change cause? Early analyses like Mendelsohn et al. (1994) estimated damages by regressing economic outcomes on measures of the climate (namely, long-run averages of weather). Concerns about omitted variable bias in these cross sectional estimates motivated a large literature that instead uses high-frequency weather fluctuations and flexibly controls for low-frequency confounders. This method has proven popular.[1] But concerns have, in turn, been raised that in removing omitted variables, these estimates also remove adaptation behavior that is important for understanding the long-run effect of climate change.[2]

This paper makes three central contributions to the estimation of climate change damages. First, in a simple production model that features weather effects on productivity, it shows that the estimation strategies based on the effect of high-frequency weather do not fully remove an important class of adaptation mechanisms: choices agents make in advance of weather shocks. The amount of adaptation left in the estimate depends on the degree of ex ante adjustment by agents and the autocorrelation between information about upcoming weather and the weather itself. This makes it challenging to incorporate the resulting damage estimates into a policy framework or integrated assessment model because one does not know, a priori, how much the estimates capture the damage holding adaptation fixed versus the damage net of adaptation.

Second, the paper shows that this identification issue can be solved by adding weather forecasts to the estimating equation. Weather forecasts help isolate shocks or surprises in the weather. The resulting estimates are thus purged of forward-looking adaptation and provide a clearly defined starting point for modeling impacts of climate change.[3] By using shocks to identify damages, the resulting estimates are also less susceptible to potential high-frequency confounders.

Third, including forecasts in the estimating equation provides the helpful additional feature of generating estimates of the marginal benefits of adaptation. Marginal

[1]The literature can be traced back to Schlenker et al. (2005), with important developments of the alternative, high-frequency approach in Deschênes and Greenstone (2007) and Schlenker and Roberts (2009). Dell et al. (2014), Auffhammer (2018), Kolstad and Moore (2020) provide reviews.

[2]This concern was raised in early work on the method (e.g. Deschênes and Greenstone, 2007) and has been discussed in subsequent reviews including Carleton and Hsiang (2016).

[3]Estimates of both short-run, adaptation-free damages and long-run, net-of-adaptation damages would provide helpful calibration targets for integrated assessment models. The motivation for the popularity of the high-frequency, weather-based approach to estimating damages, however, is that credible long-run estimates are challenging to generate. Recently, Lemoine (2021) has provided a complementary approach to the one described in this paper that seeks to build up long-run damage estimates using the effect of forecasts and realizations of weather.

benefits of adaptation are identified by how firms use the forecasts to alter their input choices and how those choices translate into changes in revenue for the firm. The estimates are identified from high-frequency data and thus have stronger identification than widely-used alternative approaches to estimating adaptation benefits (discussed below) that rely on low-frequency or cross sectional variation.

The empirical strategy implied by the model is used to estimate the effect of climate variation from El Niño/Southern Oscillation (ENSO) on albacore tuna harvesters in the North Pacific. The empirical setting is particularly well-suited to illustrate the benefits of the estimation framework. First, ENSO—periodic but stochastic warming and cooling of the equatorial Pacific Ocean—is a main source of medium-term global climate variation.[4] Productivity in the albacore fishery in particular has historically been strongly affected by ENSO. Second, the fishery matches the simple model well. The fishery is healthy and not heavily restricted by policy. Firms produce a single output that is storable and tradable, allowing for analysis of either output or revenue effects.[5] Third, high-quality, high-frequency, firm-level data on the fishery is available for 35 years. Fourth, during the period spanned by these data, ENSO forecasting technology experienced a sudden paradigm shift. ENSO was thought to be unforecastable until the mid 1980s. By 1989, however, skillful ENSO forecasts were created and the National Oceanic and Atmospheric Administration (NOAA) began a program to disseminate the forecasts to firms. The paper uses unique, newly digitized data on the full history of these forecasts to study their effect on the industry.

Estimates highlight the three main points of the paper. ENSO variation strongly affects the fishery, but the magnitude of this effect is sensitive to the inclusion of ENSO forecasts. Without ENSO forecast controls, the effect of ENSO realizations on firm output is overstated by more than 40%.[6] The information in ENSO forecasts

[4]The European Centre for Medium Range Weather Forecasting states that "The ENSO cycle is the largest known source of year-to-year climate variability" (Stockdale, 2021). Glantz and Glantz (2001) calls ENSO "the second most important climate process after the changing seasons."

[5]This reduces potential output price changes in response to ENSO. The primary inputs are fuel and labor. Although other global commodity prices have been found to vary with ENSO, oil is not one of them (Brunner, 2002). Wages for fishing labor do change in response to ENSO forecasts, so in an extension to the model, I show that this does not change the identification or interpretation of the results.

[6]In other words, in this setting adaptation increases the effect of ENSO on revenue. Commonly, it is stated that removing adaptation estimates will lead to the reverse result—i.e. that direct damage estimates purged of adaptation will *overestimate* the damage from climate change. For example, Deschênes and Greenstone (2007): "The primary limitation to this approach [of using high-frequency variation] is that farmers cannot implement the full range of adaptations in response to a single year's weather realization. Consequently, its estimates may overstate the damage associated with climate change." Here the opposite is true because, as the mechanism analysis will show, firms are adapting mainly by reducing costs. This highlights an important general point that latent adaptation in

is also important to the fishery. Forecasts have a three times larger effect on output than do realizations of ENSO. Interpreting this through the lens of the model, the estimates suggest that the marginal benefit of forward-looking adaptation is large relative to the direct effect of ENSO that occurs conditional on that adaptation.

Exploiting the richness of the data, secondary results examine the mechanisms the vessels use to adapt. Overall, vessels respond to forecasts by reducing their fishing effort and expenditures during periods when adverse conditions are expected. On the intensive margin, in anticipation of ENSO, harvesters fish fewer hours per days, move less during fishing trips, and employ fewer fishing lines (a labor proxy). Similarly, within a month that the vessel chooses to go fishing, vessels fish for fewer days and take slightly fewer trips per month if they anticipate that climate conditions will be bad. Across months, harvesters avoid participating in the fishery—either by declining to enter the albacore fishery or by exiting more quickly if they are currently fishing albacore—if they expect conditions to be poor. In contrast, the effect of realized ENSO conditional on the forecasts causes little or no change in any of these behaviors. On the whole, the mechanism analysis supports the primary result: revenue falls when the forecast of ENSO is high, but firm actions are generally cost-saving. Thus, the firms insulate themselves from negative profit shocks.

Finally, the model can be extended to study firm risk tolerance and learning. I adopt the reduced form of the model from Rosenzweig and Udry (2014) to determine whether the firms in this setting are risk averse. A risk-averse firm should care both about the level of the forecast and its ex ante uncertainty. In this setting, firms do appear to be risk averse. The past accuracy of ENSO forecasts (as measured by recent, historical squared forecast error) and a narrowing of the dispersion of the members of the forecast ensemble both cause higher levels of adaptation. Second, firms with more ENSO experience are better able to adapt than novice firms.

Overall, the results show that information has enabled substantial adaptation to climate variation from ENSO in the North Pacific albacore fishery. The same empirical method is potentially widely applicable because ENSO is a major source of climate variation around the world. ENSO realizations have been shown to affect many important economic outcomes including global conflict (Hsiang et al., 2011), changes commodity prices (Brunner, 2002), agricultural productivity (Meza et al., 2008), and a wide variety of infectious diseases including malaria and cholera (Kovats et al., 2003). Forecasts of ENSO might be valuable for studying adaptation across these and other settings.

current estimates of climate damage could be biasing damages up or down.

More broadly, expectations of weather likely affect many existing estimates across multiple fields in economics. Routinely updated, short-range weather forecasts have been available for all parts of the globe for decades (Bauer et al., 2015). Seasonal forecasts, though less skillful, are also available globally on a routine basis, with particular focus on helping the agricultural sector (Meza et al., 2008, Barnston et al., 2010, Toth and Buizza, 2018). Long-run climate forecasts are being used to inform financial decisions (Schlenker and Taylor, 2021), and they have been shown to affect Ricardian estimates of climate damage in agriculture (Severen et al., 2018).[7] Even without modern forecasting technology, people have been forming expectations of weather for millennia (NASA, 2002).[8] If a researcher does not have access to forecasts to include when estimating climate and weather effects, the logic underlying the estimation strategy in this paper can still be helpful. First, confounding from omission of an expectations proxy will be worse in cases where expectations play an important role: if the agents have substantial opportunity to adapt and information on which to act (the same conditions that affect value of information). Second, expectations proxies aside from forecasts can be used to reduce confounding. Fixed effects as well as lags and leads of weather can all be useful ways to reduce expectation-based confounding. Including a forecast is useful when either: (1) expectations are time varying, (2) one wants to explicitly estimate the marginal benefit of adaptation rather than simply treating adaptation as a nuisance parameter, or (3) persistent weather effects reduce the quality of lagged realizations as pure expectation proxies.

Finally, implementing an empirical strategy that combines forecasts and realized weather allows for the reduction of potential confounding from omitted variables that comes with the use of high-frequency variation while also providing estimates of the benefits of adaptation. Other work that estimates the overall effect of adaptation generally does so by comparing responses to high- and low-frequency weather variation.[9] This class of methods has two limitations. First, the results from this paper

[7]One of the aims of the Ricardian estimation approach used by Mendelsohn et al. (1994) and others is to capture effects of climate net of all adaptation actions. Severen et al. (2018) use forecasts to make the important additional point that even this approach fails to fully capture ex ante adaptations. This current paper shares a similar goal but focused on the high-frequency damage estimation literature where the issue is not that people thought all adaptation was captured but rather that little adaption was captured.

[8]Or as Roberts (2017) poetically puts it, "Once we humans began to depend on planted crops and domesticated animals, our new mode of life absolutely required us to think ahead: to anticipate setbacks and think through solutions, to plan, to map out the future world—indeed, many potential future worlds."

[9]The intuition is that high-frequency variation in weather identifies without-adaptation effects while lower-frequency variation (including cross-sectional average weather) identifies with-adaptation effects. If so, the two can be compared to estimate the value of adaptation. Examples include Dell

show that adaptation can still affect estimates based on high-frequency variation, so the comparison of the two estimates is not guaranteed to identify the benefit of adaptation. Second, the modern, weather-based empirical literature started as a way to address omitted variable bias in earlier analyses of climate impacts (Schlenker et al., 2005). Using low-frequency variation to estimate adaptation risks reintroduction of omitted variable bias concerns.

The rest of the paper proceeds as follows: Section 2 formalizes identification issues and solutions. Section 3 gives background on the empirical setting and discusses the data. Section 4 lays out the specific empirical analysis that will be performed on the data, and Section 5 reports the results of estimating that model as well as robustness checks. Section 6 investigates adaptation mechanisms over multiple time horizons. Section 7 examines heterogeneity in the adaptation response and draws out additional implications of forecast-driven adaptation. Section 8 concludes.

2 Identifying damages in the presence of adaptation

2.1 Defining direct damage from weather and adaptation

The model serves to highlight the identification challenge when agents act on expectations of weather. All aspects of the model follow a standard profit maximization setup with one complication: the firm faces costly adjustment such that inputs needs to be chosen before the arrival of weather shocks.[10]

Consider a firm choosing inputs to produce a univariate output at time t, with productivity determined, in part, by a random weather shock.[11] At the beginning of each period, the firm's problem is to maximize expected profit

$$\max_{\mathbf{x}} \mathsf{E}_{t-1}[\pi_t] = p_t f(\mathbf{x}_t) \mathsf{E}_{t-1}[g(Z_t)] - \mathbf{c}_t' \mathbf{x}_t \tag{1}$$

Output prices are denoted by p, $\mathbf{x}$ is the J-dimensional vector of inputs,[12] $\mathbf{c}$ is the

et al. (2012), Hsiang and Narita (2012), Schlenker et al. (2013), Moore and Lobell (2014), Burke and Emerick (2016). A related method compares estimates derived from high-frequency variation across entities or characteristics as in Auffhammer (2022) and Carleton et al. (2022).

[10] Model extensions are given in Section A, including the case where inputs can be chosen after weather has realized. For an extension to the case with finite adjustment costs, see Downey et al. (2023).

[11] The maximization problem resets each period. In the framework of Lemoine (2021), this assumption allows for the identification of climate damages from weather and forecast effects because the intertemporal complementarity of actions is zero. This assumption appears reasonable in many fisheries (Costello et al., 2001). In the empirical results, outcomes are analyzed at the monthly level, and harvesters rarely take trips lasting for more than a month.

[12] I use a vector of inputs to emphasize the potentially high dimensionality of the firm's adaptation choices (and the attendant identification and measurement challenges).

J-dimensional vector of input prices,[13] and Z is a stochastic weather variable with at least one finite moment that affects revenue via the function g.[14] Further assume that $f(\mathbf{x})$ is twice continuously differentiable and concave.[15] As is standard, a subscript on an expectation operator denotes the information set on which the expectation is conditioned, so $\mathsf{E}_{t-1}[g(Z_t)] \equiv \mathsf{E}[g(Z_t)|\mathcal{F}_{t-1}]$ is the expected effect of weather this period conditional on information about the weather in all time periods up to and including the most recent period.

The problem is a standard one, as indicated by the representative first order condition,

$$p_t \mathsf{E}_{t-1}[g(Z_{it})] \frac{\partial f(\mathbf{x}_{it})}{\partial x_{jit}} = c_{jt} \tag{2}$$

which says that at optimum the firm equates the expected marginal benefit of an input change with the marginal cost.

The first order conditions make three things clear. First, adaptation is the set of changes in all inputs in response to an expected change in weather. Optimized inputs implicitly defined by Equation (2) can be denoted $x_{jt}(p, \mathbf{c}, \mathsf{E}_{t-1}[g(Z_t)])$ for all j and t, so the formal definition of adaptation is

$$\mathbf{A}_t \equiv \left(\frac{\partial x_{1t}(p, \mathbf{c}, \mathsf{E}_{t-1}[g(Z_t)])}{\partial \mathsf{E}_{t-1}[g(Z_t)]}, \ldots, \frac{\partial x_{Jt}(p, \mathbf{c}, \mathsf{E}_{t-1}[g(Z_t)])}{\partial \mathsf{E}_{t-1}[g(Z_t)]} \right)' = \frac{\partial \mathbf{x_t}}{\partial \mathsf{E}_{t-1}[g(Z_t)]} \tag{3}$$

This formalizes the idea that adaptation is the set of actions taken to help reduce the negative effects of a potential change in the environment or to capitalize on gains from such a change.[16] This formalization is helpful to generalize from the single adaptation strategy or mechanism—staying indoors on hot or polluted days (Neidell, 2009), changing the mix of crops or the use of agricultural inputs (Hornbeck and Keskin, 2014, Sloat et al., 2020), or use of air conditioning (Barreca et al., 2016)—to

[13] Section A.1 considers a model where prices are a function of climate. Input price changes do not affect the interpretation of the empirical results. Output price changes do not affect the direct damage estimate but will change adaptation estiamtes. In the empirical setting, the assumption of output prices being uncorrelated with weather and forecasts is testable and appears to hold (see Section C).

[14] The function g could, for example, capture the fact that moderate temperatures are beneficial for the firm while extreme temperatures are harmful. Multiplicative separability between weather and inputs is assumed here for ease of presentation. In the empirical setting, interactions between forecasts and realizations have a practically small effect on firm output (see Section B.8).

[15] See Section A for the extension to discontinuous inputs. Identification remains unchanged, but the welfare conclusions discussed below will change. The function g need not be differentiable because the firm is not directly choosing Z.

[16] For examples of such a definition, see EPA (2017) or IPCC (2014).

the overall effect of adaptation on an individual's welfare.

Second, in the continuous case, optimal adaptation is determined by an equivalence between the marginal cost of adapting and the marginal benefit of adapting. The return on each adaptation mechanism is a function of the marginal productivity of each input as well as the expectation of the firm about the future state. This equivalence suggests that, in principle, estimates of adaptation could come from exogenous changes in any of these variables. To estimate overall adaptation benefits or costs, however, one would need to have exogenous price variation for all adaptation mechanisms or shocks to all marginal products. Aside from the high data hurdle, such a procedure requires the researcher to know the full set of available adaptation mechanisms a priori.

Third, using expectations allows the researcher to be agnostic about the set of available mechanisms because expectations will capture the reduced form effect of all forward-looking adaptation decisions. Denote realized revenue by $y_t \equiv pf(\mathbf{x}_t)g(z_t)$ and ex ante revenue as the expectation of this term with respect to information at $t-1$. In the model, the marginal benefit of adaptation is the adaptation vector multiplied by the revenue value of those changes, denoted

$$B(\mathbf{A}_t) \equiv \frac{\partial \mathsf{E}_{t-1}[y_t]}{\partial \mathbf{x}_t} \frac{\partial \mathbf{x}_t}{\partial \mathsf{E}_{t-1}[g(Z_t)]} \tag{4}$$

where arguments of the maximized output and choice variables have been suppressed for brevity. Producing estimates of this value is one of the primary benefits of including measures of expectations in one's analysis. Understanding the benefit of adaptation is important for generating accurate estimates of the direct effect of weather, as will be described below. The benefit of adaptation is also useful for bounding adaptation costs, which need to be taken into account when assessing the benefits of policy to address environmental externalities that would save on such costs.

The direct effect of weather is the residual effect conditional on adaptation. In the context of the model, the direct effect is

$$D \equiv \frac{\partial \mathsf{E}_{t-1}[y_t]}{\partial \mathsf{E}_{t-1}[g(Z_t)]} = pf(\mathbf{x}_t). \tag{5}$$

Identification of this object is the primary goal of the paper.

2.2 Identification and confounding of damage estimates

The theory presented so far shows that if a researcher observes the beliefs agents hold about the weather and has access to ex ante data, then both the value of adaptation and the direct effect of weather can be estimated. Here, I show that these values can also be identified using ex post data and the bias that can result from failing to control for expectations.

Formally, inputs are a function of expected weather (versus realized weather), so

$$\mathsf{E}_{t-1}[f(\mathbf{x}^*(p, \mathbf{c}, \mathsf{E}_{t-1}[g(Z_t)]))] = f(\mathbf{x}^*(p, \mathbf{c}, \mathsf{E}_{t-1}[g(Z_t)]))$$

Thus, changes in realized weather identify the direct effect because

$$\frac{\partial y_t}{\partial g(z_t)} = pf(\mathbf{x}^*) = \frac{\partial \mathsf{E}_{t-1}[y_t]}{\partial \mathsf{E}_{t-1}[g(Z_t)]}$$

For identification of the adaptation effect, note first that with respect to the information at time $t-1$, $\frac{\partial \mathbf{x}^*(p,\mathbf{c},\mathsf{E}_{t-1}[g(Z_t)]))}{\partial \mathsf{E}_{t-1}[g(Z_t)]}$ is known, so $\mathsf{E}_{t-1}[\partial \mathbf{x}^*/\partial \mathsf{E}_{t-1}[g(Z_t)]] = \partial \mathbf{x}^*/\partial \mathsf{E}_{t-1}[g(Z_t)]$.

Showing that $\mathsf{E}_{t-1}[\partial y_t/\partial \mathsf{E}_{t-1}[g(Z_t)]] = \partial \mathsf{E}_{t-1}[y_t]/\partial \mathsf{E}_{t-1}[g(Z_t)]$ requires an interchange of integration and differentiation. The assumption of monotonicity of output with respect to $\mathbf{x}$ allows for the application of the dominated convergence theorem, so this interchange is valid. Together, then, these two results show that the expectation of the derivative of ex post output with respect to expected weather recovers the partial derivative of ex ante output with respect to expected weather.

These results imply that estimates of the direct effect and the marginal benefit of adaptation can be generated by regressing firm revenue on expectations ($\mathsf{E}_{t-1}[g(Z_t)]$) and realizations ($g(z_t)$) of weather:

$$y_t = \alpha_0 + \alpha_1 g(z_t) + \alpha_2 \mathsf{E}_{t-1}[g(Z_t)] + \nu_t, \tag{6}$$

Excluding the measure of expectations causes the direct effect to not be identified. In that case, the error term is $\tilde{\nu}_t \equiv \alpha_2 E_{t-1}[g(Z_t)] + \nu_t$. The bias can be derived from the usual omitted variable bias formula, which yields

$$\hat{\alpha}_1 \xrightarrow{p} \alpha_1 + \alpha_2 \frac{\mathsf{Cov}(\mathsf{E}_{t-1}[g(Z_t)], g(z_t))}{\mathsf{V}(g(z))} \tag{7}$$

where the arrow indicates a probability limit as the number of observed time periods goes to infinity and Cov is the covariance; V the variance of the indicated variables. The bias will depend, first, on how much forward-looking adaptation occurs and in which direction it affects revenue (α_2). Second, the bias will depend on the covariance between expectations and realizations of weather normalized by the variance of weather. This is precisely the ordinary least squares coefficient one would get from regressing the forecast of weather on weather itself. For a well-behaved forecast, this term will be positive and less than 1.

The direction of the bias from excluding expectations from the analysis ultimately depends on whether the direct weather effect and the adaptation benefit effect go in the same direction. If they do, then the direct effect will be over-estimated by an analysis that fails to include expectations. If the two effects go in opposite directions, then excluding expectations will lead to an under-estimate of direct damages.

2.3 Using public forecasts to measure beliefs

Given the identification argument presented above, the ideal estimating equation to measure the benefit of adaptation and the direct effect of weather is Equation (6) where $\mathsf{E}_{t\ 1}[g(Z_t)]$ is the private expectation that the agent holds about the weather next period. Finding good proxies for agent expectations is challenging in general, motivating the large literature on how to elicit beliefs or deal with them econometrically (Muth, 1961, Manski, 2004).

Researchers studying weather effects, however, are well positioned to employ a method with many good theoretical properties: using professional weather and or climate forecasts as the measure of agent expectations. Modern weather forecasts are formal statements of the expectations of the forecaster about future conditions, and many individuals and firms rely on these forecasts to make weather-contingent plans. Therefore the forecasts have the potential to capture some or all of the information contained in the expectations of private agents while also being observable.

Professional forecasts will provide a good measure of agent beliefs under the assumptions that: (1) the forecasts are public, and (2) agents are maximizing expected profit. The quality of the proxy will depend on the degree to which the forecasts capture the full information available to agents. To see this, denote the public forecast as $\widehat{g(z)}$, and consider the public forecast as a proxy for the private expectation.

The first condition for a good proxy is that it is redundant with the variable for which it proxies (Wooldridge, 2010, ch.4), meaning that if the true expectations of the agent were observed, then the public forecast would not be helpful in explaining revenue. Formally, that $\mathsf{E}\left[y\middle|g(z),\mathsf{E}_{t\ 1}[g(Z)],\widehat{g(z)}\right] = \mathsf{E}\left[y\middle|g(z),\mathsf{E}_{t\ 1}[g(Z)]\right]$. Optimization ensures that this condition will be satisfied. Private beliefs should always be either equal to or sufficient for the public forecast (if not, then the agent is losing profit by ignoring information), so conditioning on public forecasts will not add any information relative to conditioning on private forecasts.

The second condition for a forecast to be a good proxy is, informally, that it removes the endogeneity of realized weather that occurs if agent expectations are not taken into account in Equation (6). Projecting private beliefs onto public forecasts

$$\mathsf{E}_{t-1}[g(Z_t)] = \theta_0 + \theta_1 g(\hat{z}_t) + \xi_t \tag{8}$$

this condition can be formalized as saying that if the researcher regresses revenue on realized weather, $g(z)$, and the public forecast using

$$y_t = \alpha_0 + \alpha_2\theta_0 + \tilde{\alpha}_1 g(z_t) + \theta_1\alpha_2 g(\hat{z}_t) + \alpha_2\xi_t + \nu_t. \tag{9}$$

then the covariance between realized weather and the error term needs to be zero. Zero covariance between ν_t and $g(z_t)$ follows from the assumption that Equation (6) is well identified. The condition thus amounts to needing $\mathsf{E}[g(z_t)\xi_t] = 0$. Under this condition, the estimate of the direct effect, α_1, will be consistent by the usual arguments for the consistency of the ordinary least squares estimator. A sufficient condition for this to hold is that the public forecaster has a weakly larger information set than the private agent. In such a case, the agent will adopt the public forecast as their private belief. Elaboration on this case can be found in Section A.4.

The adaptation effect, α_2, can be identified under a substantially weaker assumption. To get correct inference on this parameter, the researcher only needs that θ_1 be equal to 1. A sufficient condition for this to hold is that the private and public forecasts are both unbiased estimates of $g(z_t)$. In that case, $g(\hat{z}_t)$ will be an unbiased estimate of $\mathsf{E}_{t-1}[g(Z_t)]$ as well, so $\theta_1 = 1$ and $\theta_0 = 0$. Section 3.1 provides evidence that unbiasedness is the stated goal of forecasters in the empirical setting.

An alternative approach to measuring agent expectations that is employed implicitly in much of the literature is to use average weather in the form of location fixed effects. When studying climate adaptation, using average weather might not provide good inference. First, climate change implies that the distribution of weather is shifting over time, so if agents are updating their beliefs about the climate, then historical averages will not be perfectly accurate proxies for agent beliefs.[17] Using contemporary averages also makes the assumption that agents have, and act on, perfect information about average temperature. This will lead to attenuation of adaptation estimates in cases where agent beliefs do not perfectly match realized changes in climate. This method further assumes that the period over which weather is averaged is equal to

[17]The error in this approximation can be large in extreme cases. For instance, if agents have perfect foresight and the mean of the climate process is drawn from a stochastic process with no serial correlation, then the historical average weather will have zero correlation with the expected weather this period.

the period over which beliefs about the weather are fixed.[18] Finally, forecasts vary at the same frequency as the underlying weather, allowing for estimation of adaptation benefits and direct effects with equally strong identification arguments (the same low-frequency confounders can be removed from both estimates). Recent methods that use low-frequency variation to study adaptation run the risk of conflating adaptation with confounders that are not present in their estimates of the direct effect based on high-frequency data.

2.3.1 Violations of forecast proxy conditions

In many cases where the forecast proxy conditions are violated, the adaptation estimate will be attenuated and the direct effect will be larger in magnitude—both leading to underestimates of the relative degree of adaptation. Thus, the method presented here provides a conservative estimate of adaptation under plausible assumptions. The discussion below explores the implications of a series of relaxations of the assumptions in the previous section.

First, maintain the assumptions that forecasts are public and that agents are fully sophisticated. But make no assumption about the relationship between the public and private forecasts. Then an optimizing firm's private forecast will only differ from the public forecast if there is additional predictive power in the private forecast. In that case one should expect that $\mathsf{E}[g(z_t)\xi_t] > 0$, so the omitted variable bias formula can be applied to find that $|\hat{\hat{\alpha}}_1| \xrightarrow{p} \left|\alpha_1 + \alpha_2 \frac{\mathrm{Cov}(\ ;g(z))}{\mathrm{V}(g(z))}\right|$. If $\alpha_2 < 0$, then the estimated coefficient will be biased upward, meaning that the direct effect will be over-estimated. As discussed above, an extreme version of this is omission of ***any*** measure of an agent's beliefs.

Second, perhaps due to ensemble averaging considerations following Efron and Morris (1975), a firm or the forecaster might prefer a biased estimator. If the level of bias is constant, the bias will enter θ_0, and the estimate of the adaptation effect will still be consistent for the true adaptation effect. The covariance between ξ_t and realized weather will no longer be zero, and the inconsistency will depend on the sign of the bias of the estimator employed by the forecaster or agent.

Finally, if the firm is not optimizing and creates its own forecasts with a smaller information set than the public forecaster or if the firm and forecaster information sets are partly disjoint, then one could see bias in α_2. For instance, if the firm consumes its

[18] A final issue that applies to the empirical setting of this paper is that average weather cannot be used in cases where the relevant climate shifts are measured in terms of anomalies from historical averages. The expected value of the process over any sufficiently long period in this case will be zero by construction, so no identifying variation in average weather will exist.

own forecast even though it is inferior to the public forecast, then the public forecast would possess measurement error when used in the estimating equation. In general, so long as the public forecast is positively correlated with the realized state, then unless the private agent has a reason to construct a negatively correlated forecast, including the public forecast in the estimating equation will return the correct sign on the adaptation effect. It will also help reduce the omitted variable bias from ignoring adaptation.

3 Empirical setting, background, and data

3.1 Albacore fishing, ENSO, and ENSO forecasting

The remaining sections of the paper apply the theory from Section 2 to estimate the direct effect of climate fluctuations and the marginal benefit of adaptation for firms in the U.S. North Pacific albacore tuna fishery. Four attributes of this setting make it ideal to study adaptation. First, the fishery is affected by ENSO, an important climate phenomenon that causes changes in oceanic and weather conditions (and therefore affects fishing quality). Second, for multiple decades, the fishery has relied on professional forecasts of ENSO. NOAA issues ENSO forecasts directly to albacore harvesters in the fishery, and interviews with harvesters indicate that these forecasts are utilized. The fishery is also almost entirely located in the northern Pacific Ocean, far from where ENSO conditions develop. This means that NOAA information is plausibly the primary or only source of ENSO information for these firms. Third, concerns about other confounding effects are minimal. The fishery does not suffer from congestion, is not subject to catch quotas, and the albacore population is relatively healthy (Albacore Working Group, 2014). The U.S. harvesters studied here account for a small part of the global albacore tuna output. A large portion of albacore tuna is canned and therefore storable, reducing output price effects from climate variation. A primary variable cost comes from diesel fuel, a globally traded and produced commodity.[19] Fourth, detailed logbook records of output and some inputs are legally required to be kept on a daily basis for each firm in the industry, providing more than 35 years of high-quality data.

Albacore (*Thunnus alalunga*) typically stay in waters with sea surface temperature between 15 and 20 C. They also follow oceanic fronts with strong temperature gradients which limit the movement of their prey. The temperature preferences of

[19]Empirically, output and fuel prices do not respond to changes in ENSO or ENSO forecasts during the sample period. See Tables A7 and A8. Overall fishing industry wages in the area do respond to ENSO forecasts. Section A.1 shows that direct effects and adaptation benefits are still identified in this case.

albacore make them highly responsive to changes in climate.

ENSO affects temperature in the North Pacific (see Figure A5) as well as oceanic conditions like temperature gradients. These shifts make it harder for vessels to locate albacore (Fiedler and Bernard, 1987).[20] ENSO, therefore, generally entails more intensive and costly search for fish. In interviews, harvesters indicate that they respond by temporarily exiting the albacore fishery, a response that I confirm in Section 6.3 (Wise, 2011, McGowan et al., 1998).

On average, harvesters take fishing trips that last two weeks, but trips can last up to three months. Harvesters generally take between 1 and 2 trips per month. An ideal trip involves an initial transit to a fishing ground followed by little movement of the vessel as actual fishing occurs. Because ENSO effects are felt in the fishery as quickly as a week after equatorial temperature changes (Enfield and Mestas-Nuñez, 2000), this strategy can be disrupted by unanticipated ENSO events.

Unfortunately for the harvesters, prior to the late 1980s, ENSO was not forecastable. In fact, despite the importance of ENSO to global climate, equatorial temperature anomalies were often not even ***measured*** prior to the deployment of the Tropical Atmosphere Ocean (TAO) array of weather buoys starting in 1984.[21]

Skillful forecasts of ENSO were developed starting in the mid-1980s. Cane et al. (1986), a group of researchers at Columbia University's Lamont-Doherty Earth Observatory (LDEO), published the first coupled ocean-atmosphere forecast of ENSO, named LDEO1. A stated goal of the LDEO forecasting group was to produce unbiased forecasts of ENSO (Chen et al., 2000). In the late 1980s, NOAA's Climate Prediction Center (CPC) began to produce a statistical forecast of ENSO based on Canonical Correlation Analysis (CCA).

Starting in June 1989, NOAA's National Centers for Environmental Prediction (NCEP) began publicly issuing three-month-ahead ENSO forecasts in the Climate Diagnostics Bulletin, a publication containing global climate information and medium-term climate forecasts. The Climate Diagnostics Bulletin initially reported the LDEO1 forecast, and forecasts from additional forecasting groups were incorporated as they were introduced, starting with the CCA forecast in November 1989.[22] By the end of

[20] Lehodey et al. (2003) shows that, in addition to spatial dislocation, Pacific albacore recruitment tends to fall after El Niño periods, indicating that there might be temporal spillovers between ENSO and catch in the fishery. I check this in Table 3 and rule it out as an explanation of the main results.

[21] NOAA's history of ENSO measurement notes, "Development of the Tropical Atmosphere Ocean (TAO) array was motivated by the 1982-1983 El Niño event, the strongest of the century up to that time, which was neither predicted nor detected until nearly at its peak" (NOAA, 2013).

[22] For examples of these historical Bulletins, one can see the archive going back to 1999 online (NCEP, 2020).

the sample in 2016, the Bulletin published 21 ENSO forecasts on a monthly basis. See Appendix B.1 for more information on the content of the Bulletins.

The forecasts were an improvement over simple persistence or autoregression-based predictions. Regressing the standard measure of ENSO—the Niño 3.4 index—on the forecast time series, the coefficient is 1.03 (standard error of 0.04) and the R^2 is 0.73. In contrast, regressing Niño 3.4 on its three-month lag, the coefficient is 0.74 and the R^2 is 0.55. Thus, the forecasts provide about a 25% improvement, on average, over a simple statistical forecast. Analyses of forecast accuracy and performance over time can be found in Barnston et al. (2010, 2012) as well as Figure A4, which shows the forecast skill over and above a persistence forecast. The figure shows that there has been some variation in forecast quality over time but that the forecast has been consistently skillfull since the early 1990s.

Around the same time that ENSO forecasts were being created, NOAA started a program called CoastWatch, first launched in 1987, to disseminate forecasts, satellite imagery, and other data to coastal businesses and individuals. ENSO forecasts from the Climate Diagnostics Bulletin were incorporated in the CoastWatch releases, and personal correspondence with albacore harvesters indicates that CoastWatch forecasts were routinely posted at albacore fishing ports along the Pacific coast. Even today, private companies selling weather forecasts and satellite imagery to the albacore fishery repackage the NOAA ENSO forecasts.[23]

For this paper, I focus on the effects of the three-month-ahead ENSO forecast. The use of this forecast is due in part to the history of NOAA's public forecast releases. The three-month-ahead forecast was the first one issued by NOAA and therefore has the longest history. Given the timing of ENSO effects being felt in the North Pacific and typical trip length, this forecast horizon is also likely to be relevant for fishing decisions.[24]

3.2 Dataset construction

For estimation, data on equatorial and North Pacific sea surface temperatures, ENSO forecasts, vessel-level fish catch, and relevant prices are combined. Summary statistics for the variables can be found in Table 1 and more details about dataset construction can be found in the Appendix (Section B).

ENSO is typically measured using a time series of temperature anomalies relative

[23]For instance, SeaView Fishing, a private firm used by the fishers that I spoke to, simply links to NOAA's ENSO forecast website for predictions of El Niño and La Niña (SeaView Fishing, 2021).

[24]The choice of horizon is also relatively unimportant if one is only interested in separately identifying the benefit of adaptation and direct effect. Use of multiple forecast horizons can be helpful for understanding the timing of adaptation.

to a 30-year temperature average for a region of the equatorial Pacific Ocean. NOAA's CPC publishes monthly average temperature anomalies in what is known as the Niño 3.4 region of the Pacific, a rectangular area between 120 W–170 W longitude and 5 S–5 N latitude. This study uses the anomalies calculated with respect to the 1971-2000 average. Following Trenberth (1997) and NOAA, I classify El Niño and La Niña events based on five consecutive months where the three month moving average of the Niño 3.4 index is greater than 0.5 C for El Niño or less than 0.5 C for La Niña.

Table 1: Summary Statistics

	Mean	St. Dev.	Obs.
Monthly number of fish caught	185.85	833.48	120,693
Monthly catch (tons)	0.93	4.59	120,693
Niño 3.4 index	0.07	0.88	120,693
3 month-ahead Niño 3.4 forecast	0.03	0.69	120,693
Vessel length (m)	16.68	5.78	115,095
Fuel price (2001 $/L)	0.41	0.20	120,302
Albacore price (2001 $/kg)	2.63	0.54	120,693

Notes: Averages, standard deviations and number of observations for primary variables in the dataset are shown for the estimation sample (September 1989 to December 2016, excluding January each year and observations without albacore price). Observations are at the vessel-month level.

Data on ENSO forecasts come from two sources. Public ENSO forecasts have been issued as part of NOAA's Climate Diagnostics Bulletin since June 1989. These are usually published as a time series of point forecasts for the coming few months or seasons, along with observations of ENSO from recent months. I digitized forecasts from these bulletins for the period from 1989 until 2002. In 2002, the International Research Institute for Climate and Society (IRI) began keeping records of publicly issued ENSO forecasts, and Anthony Barnston at IRI provided digital records for the period from 2002 to the present. For the analysis, I use the three-month-ahead forecasts, for reasons discussed in Section 3.1. Because I use three-month-ahead forecasts, my sample begins in August 1989 (the target date of the first operational forecast issued in June). The sample ends in December 2016. More details on the construction of the historical forecast dataset can be found in Appendix B.1.

The data for the albacore fishery consist of daily, vessel-level logbook observations of U.S. troll vessels.[25] The National Marine Fisheries Service (NMFS) requires the

[25] These records begin in 1981. My primary estimation sample begins with the introduction of forecasts in 1989, but I do some supplementary analyses on the records from 1981 until 1989 in Section B.6.

vessel operator to maintain accurate logbook records in order to access the fishery. All fishing days are observed, with additional information provided for some transiting and port days (these latter data are not consistently reported). For each fishing day, the logbooks record the number of fish caught, the weight of fish, a daily location record (latitude and longitude), the sea surface temperature, the number of hours spent fishing (versus steaming, baiting, or doing other activities), and the number of troll lines used. At the trip level, the logbooks record vessel length, departure and arrival port, and total weight of catch for the trip. Weight is observed at the daily or trip level for more than 98% of the sample. Weight is interpolated for the remaining observations. The daily location reported in the logbooks is used both for spatial clustering of standard errors, as detailed below, and to calculate distance traveled each day (by taking the great circle distance between points on consecutive dates).

Landing port is matched to the Pacific Fisheries Information Network (PacFIN) database of annual albacore sale prices (ex-vessel prices) for 1989 to 2016. Only ports in the continental U.S. are in the PacFIN database (about 78% of the primary estimation sample). I perform my primary analysis on the sample where output price is observed and show robustness of the albacore catch results to inclusion or exclusion of the remaining part of the sample. I also exclude January of every year because no fishing is ever recorded in that month over the sample period. Therefore, the primary estimation sample consists of all monthly observations of active vessels in the fishery who land fish at continental U.S. ports from February through December between August 1989 and December 2016.[26]

The vessels in the sample use #2 marine diesel fuel. Where available, the price for this fuel is used for cost calculation, but the price for this exact fuel type is not available over the full sample. From 1989 to 1999, monthly, state-level average prices for diesel, gasoline, or number 2 distillate (the class of fuel containing diesel and heating oil) are available from the Energy Information Agency "Retailers' Monthly Petroleum Product Sales Report." Different states have records for diesel fuel prices starting at different dates, but by 1995, all states in my sample report diesel prices. For periods prior to 1995 when a state does not report diesel prices, #2 distillate prices are used if they are available. Over the sample where both diesel and distillate prices are observed, the values correspond closely. If neither diesel nor distillate prices are available, then gasoline prices are used after accounting for seasonal differences between gas and diesel. From 1999 to the end of the sample, monthly, port-level prices

[26] A vessel is defined as active in the fishery for a given year if it catches albacore at any point during that year.

for marine diesel are available from the Pacific States Marine Fisheries Commission EFIN database (PSMFC, 2020). All prices have been deflated to 2001 dollars using the monthly core consumer price index from the U.S. Bureau of Labor Statistics.

Data on wages and labor in the fishing industry come from the U.S. Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW). The dataset provides quarterly, county-level measures of employment and average weekly wage for employees in California, Oregon, and Washington state over the sample period who work in firms with the 4-digit NAICS code 1141.

Finally, full costs, expenditures, and revenues for a panel of 35 albacore harvesters was recorded from 1996 to 1999 in the National Marine Fisheries Service/American Fisheries Research Foundation (NMFS/AFRF) Cost Expenditure Survey. These are the best available data for costs in this fishery, and the fraction of costs attributable to fuel and labor is calculated based on this sample.

4 Empirical strategy

The conceptual model shows that to estimate the effect of ENSO on the fishery one can regress revenue on forecasts and realizations of ENSO, as in Equation (6). In the primary results, I estimate linear specifications regressing revenue or output on the one-month lag of ENSO and the three-month ahead forecast of that ENSO realization.[27]

Linearity has important advantages for simplicity of interpretation and estimation. Semiparametric tests using pre-forecast data also support the use of a linear estimating equation. I observe logbook records starting in 1981, prior to the existence of public forecasts. Under the assumption that month-to-month changes in ENSO were unforecastable (plausible given the history of ENSO measurement and forecast development), I estimate the effect of ENSO on output in the period prior to the introduction of forecasts to provide evidence on the correct functional form without needing to further account for agent beliefs. Figure 1 implements this test, showing the semiparametric relationship between the one-month lag of ENSO conditional on baseline controls (discussed below) and output conditional on the same covariates. Importantly, these controls include additional lags of ENSO, so the figure shows the effect of month-to-month changes in ENSO. The figure shows that the relationship between ENSO and output in this period was linear.[28]

Lagged effects plus linearity imply that the effect of ENSO on the firm can be

[27]The lag allows time for changes in ENSO in the equatorial Pacific to begin affecting weather and oceanic conditions in the North Pacific.

[28]Further details on this estimation can be found in Section B.6.

Figure 1: Semiparametric Relationship Between Output and ENSO Before Public Forecasts Existed

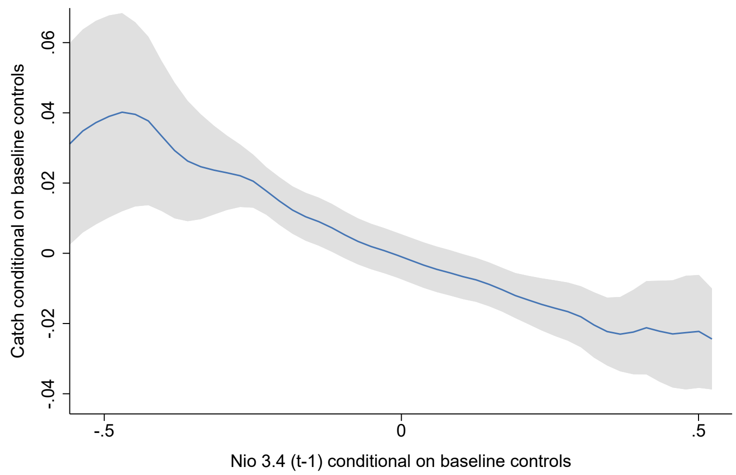

Notes: The figure shows a local linear regression (Epanechnikov kernel with bandwidth of 0.13) of monthly catch on the Niño 3.4 index the previous month. Both variables are residualized on month of year, year, and vessel fixed effects as well as two additional lags of the Niño 3.4 index. The Niño 3.4 index is Winsorized at the 1% level to improve legibility. The sample is from 1981 to May 1989 before ENSO forecasts were released. Shaded area gives the 95% confidence interval.

written

$$g(z_{t-1}) = \gamma_{\ell;0} + \gamma_{\ell;1} z_{t-1} \tag{10}$$

where $\gamma_{\ell;0}$ is a positive constant large enough to induce entry in to the fishery, z is a measure of ENSO, and $\gamma_{\ell;1}$ captures the effect of temperatures last month in the equatorial Pacific on the fishery. If warmer temperatures are harmful for productivity in the fishery, then $\gamma_{\ell;1}$ will be negative. If cooler temperatures are harmful, then $\gamma_{\ell;1}$ will be positive.

The estimating equation to identify the benefit of adaptation and direct effect of ENSO is then

$$y_{it} = \beta_1 z_{t-1} + \beta_2 \hat{z}_{t-1} + \mathbf{z}_t' \cdot \alpha_z + \hat{\mathbf{z}}_t' \cdot \alpha_{\hat{z}} + \delta_{1;i} + \delta_{2;y(t)} + \delta_{3;m(t)} + \varepsilon_{it} \tag{11}$$

where y_{it} is revenue for vessel i at time t (time is measured in months). The two primary variables of interest are z_{t-1}, the realized value of the Niño 3.4 index the previous month, and $\hat{z}_{t-1}$, the three-month ahead forecast of Niño 3.4 in month $t-1$.[29] The variable ε is assumed to be a vessel and time-varying, stochastic error term conditionally uncorrelated with Niño 3.4 or its forecast.

The baseline specification includes year fixed effects to account for overall changes in the fishery, climate regime, and forecasting system. Month fixed effects account for regular, seasonal patterns in fishery productivity and ENSO intensity. Vessel fixed effects adjust for stable characteristics of the harvesters, including features of the vessel and routine fishing grounds. The baseline specification also includes two additional lags of both the Niño 3.4 index and the three-month-ahead forecast (so $\mathbf{z}_{t-\ell}$ and $\hat{\mathbf{z}}_{t-\ell}$ contain z_{t-2}, z_{t-3}, $\hat{z}_{t-2}$, and $\hat{z}_{t-3}$) to isolate news in ENSO. Excluding these variables could allow lagged but persistent effects of ENSO realizations to confound the forecast coefficient. The residual variation in ENSO and the forecasts—about 30% of the unconditional variation—should isolate innovations in ENSO realizations and forecasts.

Adaptation is measured by the magnitude of the coefficient on $\hat{z}_{t-1}$. The larger the magnitude of β_2, the greater the adaptation because it means more of the effect of ENSO is operating through changes in actions by the agents. The effect of ENSO net of forecasts captured by β_1 reflects the direct effect that agents are unable to adapt away.

5 The effect of ENSO

5.1 Estimates of direct effect and adaptation

Table 2 shows results from implementing the primary identification strategy. Each column shows estimates of versions of Equation (11) using monthly data. The dependent variable in the first two columns is the number of fish caught per month by each vessel (a measure of output). In the third and fourth columns it is the revenue for each vessel. Given the different scales, all dependent variables are standardized to have mean 0 and standard deviation of 1 to aid interpretability across columns. The primary explanatory variables are listed in the left column, and control variables are indicated below the coefficient estimates. The standard errors in all models are spatial-temporal heteroskedasticity and autocorrelation robust. Spatial correlation in the error term is accounted for using the procedure from Conley (1999) using a uniform kernel centered around the recorded latitude and longitude of the vessel with

[29] Issued in month $t-4$. As discussed above, I use the three-month ahead forecast because it has the most complete data series and because it likely matches the decision-making horizon of the firms (see Section 3).

a radius of 30km. Autocorrelation in the errors is accounted for using 24 months of lags (Newey and West, 1987).[30]

Table 2: Effect of ENSO on Standardized Output and Revenue

	(1) Catch	(2) Catch	(3) Revenue	(4) Revenue
Niño 3.4	-0.091***	-0.063***	-0.13***	-0.11***
	(0.022)	(0.024)	(0.021)	(0.023)
$\widehat{\text{Niño}}$ 3.4		-0.19***		-0.16***
		(0.035)		(0.034)
Lagged controls	Yes	Yes	Yes	Yes
Vessel FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Unique Vessels	1,214	1,214	1,214	1,214
Observations	120,674	120,674	120,674	120,674

Notes: The table shows results from estimating equation (11) on monthly data. The dependent variable in each model is indicated at the top of the column. All dependent variables are standardized. Catch is the total number of fish caught per month. Revenue is the total ex-vessel value of catch. Additional controls are indicated at the bottom and are lagged Niño 3.4 index, lagged forecasts (Columns 2 and 4), and fixed effects for vessel, year, and month. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 2 year lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Overall, the table shows that realizations and forecasts of ENSO have strong, negative effects on output and revenue in the fishery. The effects of these variables are similar when using either output or revenue as the outcome. Due to the presence of some interpolation in the revenue measure, I focus primarily on output throughout the rest of the paper.[31]

The first and third columns regress catch and revenue on measures of the realized

[30] The results are also robust to using vessel or month clustering. See Section 5.2.

[31] Table A7 shows that a 1 standard deviation increase in the ENSO index is associated with a 1% increase in the wholesale price of albacore and a 1 standard deviation increase in the 1-month-ahead ENSO forecast is associated with a 3.7% decrease in price, both insignificant. Using output has the advantage of removing any confounding effect of price even if it is small.

strength of the one-month lag of the Niño 3.4 index but do not include forecasts. The results are included to illustrate the omitted variable bias that occurs if forecasts are not in the regression. The coefficients on realized Niño 3.4 indicate that ENSO has a negative effect on catch and revenue. An increase in the Niño 3.4 index from 0 to 1 (moving from normal conditions to a moderately strong El Niño) reduces catch and revenue by about 0.1 standard deviations. Translated into percent changes, a moderate El Niño leads to a 41% drop in output and a 67% drop in revenue.

Without including forecasts, however, these result do not give a complete or accurate picture of the effect of ENSO on the fishery. Columns 2 and 4 add the first lag of the three-month-ahead forecast of ENSO (row $\widehat{Niño\ 3.4}$). The two coefficients in the table correspond to β_1 and β_2 from Equation (11). One can see that predicted changes in ENSO have a much larger effect on output and revenue than do realized changes. An increase in the forecast leads to a drop in output roughly three times larger than a comparable change in realized ENSO. The effect of forecasts on revenue is roughly 1.5 times larger than the effect of realized ENSO. In percentage terms, a forecast of a moderate El Niño event leads to a 85% drop in output and 83% drop in revenue.

The first coefficients from columns 2 and 4 show that conditional on forecasts, the effect of realizations of ENSO is also reduced relative to the regressions that omit forecasts. The effect on output is overstated by 50% when forecasts are omitted. Comparing the estimates from column 1 and 2, the effect of a realized moderate El Niño event changes from a loss in output of 41% to a loss of 28%.

This drop in the effect of realized ENSO illustrates one of the main biases in climate damage estimation that can result from ignoring expectations. In the context of the model, the effect of realized Niño 3.4 identifies the direct effect of climate variation—the effect holding adaptation fixed. Omitting forecasts would lead to substantial over estimation of the direct effect in this case.

Continuing to interpret the estimates within the model, the effect of forecasts identifies the marginal benefit of forward-looking adaptation. The large forecast coefficients indicate that forward-looking adaptation is an important driver of the effect of climate variation on the industry. As will be shown in the mechanism analysis section below, the forecasts have such a large effect on production because the firms have many methods for adapting to forecasted climate fluctuations before they arrive, but, importantly, the estimator here can arrive at estimates of the overall marginal benefit of adaptation without specifying the particular adaptation mechanisms.

Why is the benefit of adaptation negative in this case? As will also be seen in the mechanism analysis, firms adapt by reducing their costs of production during periods with bad ENSO conditions. Because costs are saved, both the costs of adaptation and the benefits of adaptation are negative. In other words, costs are reduced as firms adapt to ENSO. Thus, these estimates highlight the concern about adaptation biasing the direct effect of weather but show that the bias can cause direct effects to be over or under-estimated.[32]

A second source of bias is also apparent when comparing the results with and without forecasts. The total effect of ENSO—the direct effect plus the benefit of adaptation—is underestimated by more than 50%. Summing the effects from both realized and forecasted ENSO, the estimates show that moving from normal conditions to a moderate El Niño leads to a 0.25 standard deviation decline in output and a 0.27 standard deviation decline in revenue, on average, for a vessel. If adaptation is costly, the large total effect has bearing on welfare analysis. Reducing the need for adaptation would reduce costs for firms.

5.2 Robustness

The results presented in the previous section are robust to many changes in specification and estimation strategy. Here, robustness checks are reported, with further checks shown in the appendix.

Table 3 checks the robustness of the main estimates (Table 2 Column 2) to changes in controls. In Column 1 the separate vessel and year fixed effects are replaced by a set of vessel-year fixed effects. In Column 2 the vessel and month fixed effects are replaced by vessel-month fixed effects. These more flexible controls do not appreciably change inference. Column 3 adds vessel-specific linear trends. Trends could be important because catch is rising, on average, over time, and forecast quality is also changing over time (Appendix Figure A4). Again, however trends have a negligible effect on inference.

Lehodey et al. (2003) raises the possibility that ENSO in one year might cause a drop in recruitment of fish into the harvestable stock in the next year. Column 4 shows that controlling for the level of the Niño 3.4 index from a year prior to the current month, however, does not indicate that conditions a year ago have strong bearing on adaptation to changes in ENSO this year.

Finally, Column 5 shows robustness to including additional lags of both Niño 3.4

[32]Another way to interpret the coefficient is to think about reductions in the Niño 3.4 index. These better climate conditions will lead to improved output and revenue at the expense of higher costs of production. In this case, adaptation is taking advantage of improved climate rather than buffering the firms from a worsened climate.

Table 3: Robustness to Covariates

	(1) Vessel by year FEs	(2) Vessel by month FEs	(3) Vessel trends	(4) Nino 3.4 t 12	(5) 6 lags Nino 3.4
Niño 3.4	-0.062***	-0.066***	-0.064***	-0.077***	-0.10***
	(0.021)	(0.019)	(0.024)	(0.025)	(0.026)
Niño 3.4	-0.19***	-0.18***	-0.19***	-0.17***	-0.17***
	(0.032)	(0.028)	(0.035)	(0.041)	(0.037)
SEs	Spatial	Spatial	Spatial	Spatial	Spatial
Observations	120,301	118,919	120,674	112,908	118,982

Notes: The table shows results from estimating versions of equation (11) on monthly data. The dependent variable in each model is standardized monthly number of fish caught. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects as well as two additional lags of realizations and forecasts of the Niño 3.4 index. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987), unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

and forecasts. Accounting for serial correlation is important for isolating news from the forecasts and ensuring that estimates are not polluted by persistent effects of past realizations of ENSO. Column 5, which includes 6 lags of both measures shows that changing the lag length does not alter the baseline results.

Table 4 shows variations in the standard error calculation method, changes in sample, and one additional variation in specification. Column 1 clusters standard errors at the vessel level to allow for arbitrary time series autocorrelation. The baseline estimates use Newey-West standard errors that account for correlation in the errors out to two years, so vessel clustering could be important if the degree of autocorrelation is large. One can see that the estimates are extremely precise in this case, indicating that this is not a concern.

Column 2 clusters at the year-month level. ENSO is a group shock, and forecasts are released each month, so this level of clustering allows for arbitrary cross sectional correlation in the response to those shocks. Inference is less precise in this case, but the forecast effect is still significant at the 5% level. A related, but unreported, robustness check shows that estimates are largely the same if the data are collapsed to the monthly level.

Column 3 excludes observations near Canadian fishing grounds. Congestion in the

Table 4: Robustness to Sample, Clustering, and Specification Changes

	(1) Catch	(2) Catch	(3) Catch	(4) Catch	(5) Catch
Niño 3.4	-0.063***	-0.063	-0.053**	-0.064**	-0.024
	(0.013)	(0.048)	(0.024)	(0.027)	(0.021)
Niño 3.4	-0.19***	-0.19**	-0.19***	-0.26***	-0.20***
	(0.017)	(0.095)	(0.035)	(0.042)	(0.031)
Catch t 1					0.48***
					(0.016)
Covariates	Baseline	Baseline	Baseline	Baseline	Baseline
SEs	Vessel cluster	Y-M cluster	Spatial	Spatial	Spatial
Sample	Baseline	Baseline	Latitude < 46	Drop 1997-2001	Baseline
Observations	120,674	120,674	118,923	91,527	120,674

Notes: The table shows results from estimating versions of equation (11) on monthly data. The dependent variable in each model is standardized monthly number of fish caught. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects as well as two additional lags of realizations and forecasts of the Niño 3.4 index unless otherwise noted. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987), unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

fishery is, in general, low. The exception commonly noted during interviews was due to Canadian vessels near the northern edge of the fishery. Excluding this area has a negligible effect on the estimates. Column 4 drops the period in the late 1990s and early 2000s with a historically large El Niño event. The results are largely unchanged whether including or excluding this period. Another large ENSO event occurred at the end of the sample, and excluding this period also has little effect.

Columns 5 adds the one-month lag of catch. Monthly autocorrelation in catch might be important to control for the effect of past actions. Including this control does not appreciably change the adaptation effect. The direct effect falls slightly, raising the relative importance of adaptation in determining the total effect of ENSO.

Next, some interpolation was performed to arrive at the revenue observations. Not all observations contain records of the weight of fish caught that day. For those observations, I impute weight in one of two ways. First, if the logbook records the total weight of fish caught during the trip, I multiply the number of fish caught that day by the average weight of fish for the trip. If trip-level weight is missing, then

I interpolate weight based on catch of other vessels fishing at the same time as the missing observation. Table A1 investigates whether this interpolation procedure leads to bias in the estimates. Overall, the results show that the interpolation procedure is not leading to substantive changes in estimates, in part because only about 2,000 observations are interpolated.

Output and input price effects are investigated in Tables A7 and A8. The results show that output prices and fuel prices exhibit minimal variation in response to ENSO of ENSO forecasts. Wages in the fishing sector in California, Oregon, and Washington do decline when ENSO is forecasted to be stronger. As Section A.1 shows, the effect of this input price change will be captured by the estimated benefit from adaptation.

6 Adaptation mechanisms

This section explores how the firms achieve the high rates of adaptation estimated in the previous section. From the main results, it is clear that the mechanisms are likely to be cost-saving. Firms suffer output and revenue losses due to the forecasts, so they must be saving on cost by engaging in behaviors to make the output and revenue loss worthwhile. The results below show that firms do indeed engage in multiple cost-saving measures both on the intensive margin—after choosing to go out and fish—and on the extensive margin when choosing whether to take a fishing trip in a given month. The results below are not necessarily exhaustive of all the mechanisms these firms have employed to adapt. Instead, the results are corroborating evidence that firms are primarily adapting by saving costs and reducing their exposure to downside risks.

6.1 Daily adaptation mechanisms

Table 5 shows estimates for the effect of anticipated and unanticipated changes in ENSO on choices made while fishing. The outcomes listed in the table are primarily determined on a daily or trip-level frequency. Overall, the results show that if a captain chooses to go fishing when they anticipate worse conditions, they take a variety of actions during the trip to reduce costs. If the poor conditions are unanticipated, the firm, if anything, engages in slightly more costly behavior.

The dependent variable in column 1 is hours of fishing per day. Good or bad fishing conditions could lead to more hours of fishing. If the vessel's hold was filled quickly, then fishing hours would go down. If fishing was poor, the crew might continue to fish longer to make up for the shortfall or might stop fishing earlier to change fishing locations. In response to anticipated ENSO, harvesters decrease their hours fished per day by just over 5%, a reduction in intensive-margin effort.

Table 5: Intensive-Margin Mechanisms

	(1) Hours per day Fishing	(2) Fishing lines	(3) Movement extensive	(4) Movement intensive
Niño 3.4	0.12	0.20	-18.3	55.5
	(0.23)	(0.17)	(12.7)	(53.5)
$\widehat{\text{Niño}}$ 3.4	-0.59	-1.07***	-145.3***	-396.3***
	(0.39)	(0.32)	(21.7)	(94.2)
Dep. var. mean	12.16	10.60	185.93	1,045.85
Baseline FE	Yes	Yes	Yes	Yes
Observations	12,949	15,893	120,674	15,938

Notes: The table shows results from estimating versions of equation (11) on monthly data. The dependent variable in each model is indicated at the top of each column. Additional controls are indicated at the bottom and are fixed effects for vessel, year, and month. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

One of the primary variable costs in this industry is labor, and one of the important shortcomings of the logbook data is that labor on the vessel is not recorded. The best available proxy measure is the number of fishing lines used each day, although a single worker can use multiple lines. Harvesters in this dataset use pole and line fishing—a relatively labor intensive but sustainable method where individual fishing lines are used to catch each fish. Operating more lines in this fishery requires either more effort or more labor. The results show that about 1 fewer fishing line is used per day (a 10% reduction) if ENSO is anticipated than if it is not.

Another major source of variable cost is the burning of fuel during transit and fishing. Table 5 Column 3 shows the effect of ENSO on day-to-day vessel movement. The results indicate that harvesters move less if they anticipate worse ENSO conditions. In fact, expecting a moderate ENSO event causes the harvesters to reduce movement by 80% of the average monthly movement. As will be shown below, this large effect is partly driven by the decision of whether to enter the fishery in a given month. Column 4 shows that even conditional on this decision, harvesters still move less if they anticipate bad conditions. In contrast, if bad conditions arrive unexpectedly, they move more, perhaps to compensate for the worse fishing.

6.2 Trip-level adaptation mechanisms

Many of the adaptations available to albacore harvesters can only be implemented between trips. In the extreme case, things like characteristics of the boat hull are fixed once a trip has begun. Labor is determined between trips as well, although that labor can be employed more or less intensively during the trip. Hull length of active vessels (unsurprisingly) does not change in response to ENSO. One adaptation that is available to the harvesters on a trip-level frequency and does appear to change with ENSO is the length of the trip and the number of overall fishing days in a month, as shown in Table 6.

Table 6: Mechanisms: Trip Length and Frequency

	(1) Fishing days	(2) Transiting days	(3) Trips per month
Niño 3.4	0.35	0.026	0.036
	(0.39)	(0.088)	(0.032)
$\widehat{\text{Niño}}$ 3.4	-3.31***	0.083	-0.12**
	(0.65)	(0.17)	(0.051)
Dep. var. mean	10.3	0.84	1.44
Baseline FE	Yes	Yes	Yes
Observations	15,938	15,938	15,938

Notes: The table shows results from estimating versions of equation (11) on monthly data. The dependent variable in each model is indicated at the top of each column. Additional controls are indicated at the bottom and are fixed effects for vessel, year, and month as well as two additional lags of realizations and forecasts of the Niño 3.4 index. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Column 1 shows that vessels fish fewer days per month given an expected change in ENSO. The magnitude is substantially larger than the effect for realizations of ENSO. As far as can be discerned from the data, there does not seem to be an effect of ENSO on transiting days, which are days away from port without any reported fishing. Transiting is not always reporting in the logbook records, however, so the results should be interpreted with caution. Finally, Column 3 shows that trips per month also slightly fall when a higher Niño 3.4 index is anticipated. Harvesters take about 1.4 trips per month, and they take about 8% fewer trips if they anticipate adverse conditions. In contrast, there is a small and insignificant increase in trips per month in response to realization of ENSO.

6.3 Entry and exit across months

Table 7 investigates the decision of whether or not to go fishing at all in a given month. The dependent variables in these models are short-run measures of entry and exit. *Fish this month* is an indicator equal to one if the vessel is both in the fishery and actively engaged in fishing for albacore. *Exit if shing* is equal to 1 the month a vessel exits the fishery after having fished the previous month and is 0 otherwise. The estimates are from linear probability models with spatial HAC robust standard errors and all baseline covariates. Fixed effects logit models give similar estimates for the effect of forecasts, but show no significant effects from realizations of ENSO.

Table 7: Mechanisms: Entry and Exit

	(1) Fish this month	(2) Exit if fishing
Niño 3.4	0.0061	-0.024***
	(0.0046)	(0.0028)
Niño 3.4	-0.060***	0.029***
	(0.0085)	(0.0062)
Baseline controls	Yes	Yes
Observations	120,674	120,674

Notes: The table shows results from estimating versions of equation (11) on monthly data. The dependent variable in each model is indicated at the top of the column. Additional controls are indicated at the bottom and are fixed effects for vessel, year, and month as well as two additional lags of realizations and forecasts of the Niño 3.4 index. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

The entry results show that vessels are less likely to be active in the fishery if ENSO is forecasted to be worse. This result helps explain the drop in output associated with increases in ENSO and also bolsters the movement results which indicated that some of the movement cost avoidance was done simply by not entering the fishery in a given month. Realized changes in ENSO conditional on forecasts do not have the same effect, with a precisely estimated zero effect from realizations of ENSO on the entry decision.

The short-run exit decision is not as strongly related to ENSO forecasts. This is consistent with interviews with fishers indicating that on a normal fishing trip, a captain will try to continue fishing in order to fill the hold even if the fishing is going poorly. This type of behavior might make exit less responsive to climate shocks. One does see that vessels are slightly more likely to exit if they anticipate bad conditions—again saving on costs—and slightly less likely to exit if the bad conditions are unanticipated—possibly because they need to stay out longer to fill their hold.

6.4 Net Revenue

Regressing forecasts and realizations of ENSO on output and revenue is useful for recovering both adaptation and direct effects. If all adaptation measures are continuous, then the envelope theorem says that on the margin, the benefit of adaptation will be equal to the cost of adaptation. In such a case, estimates of the marginal benefit of adaptation thus also provide estimates of the marginal costs of adaptation.[33] Estimates using profit as the dependent variable will return the direct effect of weather but not an explicit measure of the marginal benefits of adaptation.

One consequence of the profit-neutrality of intensive margin adaptation is that the effect of forecasts on profit should be zero. The logbook data do not provide details on many of the inputs necessary to calculate full profit measures in this empirical setting. In particular, there are no measures of vessel maintenance, capital costs, or the wages paid to crew. The one input that can be consistently calculated is movement during fishing trips. To arrive at movement costs, I multiply day-to-day location changes by the average real price of fuel at ports. Vessel engine characteristics are unavailable, but for vessels with known length, the average fuel consumption per kilometer conditional on vessel size is calculated from the NMFS/AFRF Cost Expenditure Survey and used to scale the fuel consumption. Fuel consumption for all other vessels is based on the unconditional average rate. The Cost Expenditure Survey shows that fuel costs represent 20% of the variable cost of running an albacore vessel, so the resulting costs are scaled to constitute that percentage of observed revenues on average.

Table 8 compares the effect of forecasted and realized ENSO on fuel costs, revenue, and revenue net of movement costs for the estimation sample where net revenue and fuel costs are observed. As expected from the movement results in Section 6.1, fuel

[33]Of course, under these assumptions, the total benefits of adaptation can still be larger than the total costs. In cases where a portion of the adaptation mechanisms are discrete, the marginal benefits of adaptation can be substantially larger than the marginal costs (Guo and Costello, 2013).

Table 8: Effect of ENSO on Net Revenue

	(1) Fuel cost	(2) Revenue	(3) Net revenue
Niño 3.4	-0.028	-0.11***	-0.11***
	(0.020)	(0.023)	(0.023)
$\widehat{\text{Niño}}$ 3.4	-0.23***	-0.16***	-0.10***
	(0.034)	(0.034)	(0.035)
Baseline controls	Yes	Yes	Yes
Observations	120,674	120,674	120,674

Notes: The table shows results from estimating equation (11) using monthly data. The dependent variable is standardized fuel cost in Column 1, standardized monthly total revenue in Column 2, and standardized revenue net of movement costs in Column 3. Additional controls are indicated at the bottom and are fixed effects for vessel, year, and month as well as two additional lags of realized and forecasted Niño 3.4 index. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

costs decline substantially when ENSO is anticipated.[34] Column 2 reproduces the estimates from Table 2 for ease of interpretation. Column 3 shows that, as predicted, the magnitude of the effect of forecasted ENSO on net revenue is smaller than the effect on revenue.[35] The effect of realized ENSO is the same for both variables. These results are in line with the theory—the direct effect is well estimated either when using net revenue or revenue, so long as adaptation is appropriately accounted for.

7 Learning and risk

7.1 Risk aversion

The theoretical model assumes that firms are solely maximizing profit. For many settings, including small-scale firms like fishing vessels, risk aversion by the vessel owner might also play an important role in decision making under uncertainty. Rosenzweig

[34] This effect is due to changes in firm behavior rather than through changes in fuel prices. Changes in ENSO do not have a substantial or significant effect on albacore or fuel prices, as shown in Tables A7 and A8.

[35] The effect falls by even more than 20%, consistent with the choice of how much to sail and expend fuel being a particularly important ex ante adaptation.

and Udry (2014) use forecasts of monsoon rain in India to investigate risk aversion in agriculture and the value of weather insurance. Adopting the reduced form of the estimating equation from that paper allows for a test of risk aversion in this setting. The expanded estimating equation becomes

$$y_{it} = \varphi_1 z_{t-1} + \varphi_2 \hat{z}_{t-1} + \varphi_3 \text{qual}_{t-1} + \varphi_4 \hat{z}_{t-1} \text{qual}_{t-1} + \mathbf{x}'_{it} \xi + \varepsilon_{RA,it} \tag{12}$$

where all variables are the same as in Equation (11) except for the new variable *qual* that is a measure of the ex ante quality of the forecast and that all of the baseline controls have been denoted by $\mathbf{x}$.

The intuition for this estimating equation is that the quality of the forecast matters for a risk averse agent when he or she is making input decisions because the quality measures how much uncertainty the forecast resolves. If the agent is risk averse, the quality of the forecast will be a moderating variable for the effect of the forecast on output. Under the maintained assumption that forecasts only affect inputs, this leads to a modification of the baseline estimating equation where forecast quality is interacted with the forecast terms.

I measure ex ante forecast quality in two ways. First, I calculate the average skill from the prior 6 months. Skill is the exponential of the log of 0.5 times the squared error of the three-month-ahead forecast divided by the squared error of a persistence forecast. See Figure A4 for the time series evolution of monthly skill. This measure is a version of the Brier skill score (Hamill and Juras, 2006) modified in two ways: first, a value of 0.5 indicates equal accuracy between a simple persistence forecast and the actual forecast. Second, all values of skill lie between 0 and 1. A value of this measure at 1 means that the forecast is perfectly accurate. Numbers below 0.5 mean that the forecast is inaccurate relative to a persistence forecast.

Theory predicts that a risk-averse agent will adapt more if skill is higher. The results in Table 9 Column 1 show that risk preferences are a potentially important factor. Harvesters adapt substantially more when skill is higher. The interaction between forecasts and skill is negative, so the benefit of adaptation is larger as skill goes up.

The second measure of quality is the standard deviation of the forecast plume in the prior 6 months (*Ensemble sq. error*). Because multiple forecasts are issued beginning in the 1990s, the standard deviation of the plume gives a summary measure of disagreement across the different forecasters. This measure is model-dependent and influenced by model errors, so it does not necessarily represent the full probability

Table 9: Assessing Risk Aversion

	(1) Catch	(2) Catch
Niño 3.4	-0.052**	-0.050**
	(0.024)	(0.023)
$\widehat{\text{Niño}}$ 3.4	-0.043	-0.23***
	(0.047)	(0.037)
Skill	-0.15***	
	(0.039)	
Skill $\widehat{\text{Niño}}$ 3.4	-0.25***	
	(0.064)	
Ensemble sq. error		-0.19***
		(0.026)
Ensemble sq. error $\widehat{\text{Niño}}$ 3.4		0.091***
		(0.015)
Baseline controls	Yes	Yes
Observations	118,982	120,674

Notes: The table shows results from estimating equation (12) on monthly data. The dependent variable in each model is standardized total catch per month. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects as well as two additional lags of realizations and forecasts of the Niño 3.4 index. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

distribution of a single forecast, but it plausibly affects the confidence that harvesters have in the projections. One should expect that a risk-averse agent will adapt less if this standard deviation measure is higher. Indeed, Table 9 Column 2 shows that if the forecast plume is wider, adaptation falls. The results also show that agents are responding to forecast-specific characteristics, lending support to the assumption that agents are directly consuming these predictions rather than reacting to something else that is simply correlated with forecast values.

7.2 Learning about ENSO and forecasts

Given the long times series available for each vessel, one can also assess the role that experience plays in forward-looking adaptation. A captain or vessel owner with more experience receiving ENSO forecasts and fishing during ENSO conditions might be better equipped to handle the adverse climate, leading to increased adaptation. On the other hand, if the forecasts turned out to be unhelpful, a more experienced captain might engage in more ex post adaptation, lowering the effect of forecasts.

Table 10: Experience with ENSO Events

		(1) ENSO	(2) El Niño	(3) La Niña
Niño 3.4		-0.16***	-0.11***	-0.18***
		(0.030)	(0.028)	(0.032)
$\widehat{\text{Niño}}$ 3.4		-0.17***	-0.19***	-0.15***
		(0.034)	(0.037)	(0.033)
Niño 3.4	Experience	0.018***	0.030***	0.032***
		(0.0031)	(0.0066)	(0.0054)
$\widehat{\text{Niño}}$ 3.4	Experience	-0.0088***	-0.0050	-0.022***
		(0.0033)	(0.0070)	(0.0060)
Forecast effect relative to total effect				
Low experience		0.61***	0.70***	0.55***
		(0.076)	(0.083)	(0.074)
High experience		0.94***	0.90***	0.93***
		(0.095)	(0.100)	(0.095)
Baseline controls		Yes	Yes	Yes
Experience trend		Yes	Yes	Yes
Observations		120,674	120,674	120,674

Notes: The table shows results from estimating a modified version of equation (11) on monthly data. The dependent variable in each model is standardized total catch per month. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects as well as two additional lags of realizations and forecasts of the Niño 3.4 index. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 10 investigates this hypothesis by including vessel-specific trends that increment each time a vessel experiences any ENSO event (Column 1), just an El Niño event (Column 2), or just a La Niña event (Column 3). Overall, the results suggest that there is an important learning effect. Vessels that have been through more ENSO events adapt at a higher rate. This relationship is summarized in the middle section of the table which shows the fraction of the ENSO effect that is due to forecasts (forecast coefficients divided by sum of forecast and realization coefficients) for novice versus highly experienced vessels. For a novice vessel (25^{th} percentile experience), the relative effect of forecasts versus realizations is about 20% lower than for an experienced vessel (75^{th} percentile experience).

8 Conclusion

Environmental impacts from a variety of source are currently large and, for many important cases, are not being addressed by collective action at a scale commiserate with the potential damages. If public policy is not appropriately aggressive, then individual adaptation will need to play an outsize role in damage reduction. Adaptation does not occur in a vacuum, however. Individuals can benefit from knowing about their own risks to make informed choices over potential adaptive responses. If individuals are already adapting based on ex ante information, this also has important implications for how estimates of environmental damage should be identified. Without accounting for adaptation, these estimates will contain some amount of latent adaptation, making it challenging to use them to assess optimal policy.

In the setting of one of the largest drivers of global climate—ENSO—and firms with flexible production, this paper assesses the impact of climate shocks and the degree of forward-looking adaptation using an estimating equation informed by a simple model of adaptation to a stochastic weather process affecting productivity. Detailed panel data and a unique set of real-time historical ENSO forecasts allow for estimation of the role of information in climate adaptation, showing that anticipation of ENSO helps harvesters take actions that substantially reduce the direct effect of this climate variable. The results also show that failing to take forecasts into account leads to a substantially biased view of the effect of realized climate shocks.

From a methodological standpoint, the empirical strategy has the potential to be applied to many settings. The novel collection of ENSO forecasts assembled for the project should allow for investigation of adaptation to ENSO processes in a number of different settings. Public forecasts of other weather, climate, and pollution processes can similarly be harnessed to understand expectation-driven behavior and arrive at cleanly identified effects of realized shocks.

Whether these estimates should influence broader discussions of optimal climate change mitigation policy hinges on extrapolating the results dynamically and across other firms. The magnitude of the change in temperature caused by ENSO—2 to 4 C for a complete El Niño to La Niña cycle—is comparable to the average warming currently being forecast for the coming century. Perhaps the more important difference when extrapolating the effects of ENSO to the effects from global climate change is that ENSO-driven changes are temporary, rarely lasting for more than two years. Therefore, attention to dynamics is critical to understanding whether the estimates presented in this paper have bearing on the effects of long-run climate change.

At least three arguments suggest that short-run adaptation estimates provide lower bounds for long-run adaptation. First, if an adaptation mechanism is inexhaustible and it is available in the short run, then it will be available in the long run. Second, if a firm owner expects a change in the environment to be permanent, then he or she will be more willing to take adaptive actions that require long-term investments. Third, technical change might improve the adaptive capacity of a given production process.

On the other hand, if adaptation mechanisms are exhausted, if agents hit corner solutions, if the prices of adaptation mechanisms rise too rapidly, or if climate change causes more extreme weather impacts, then short-run adaptation estimates will not be as good of a guide for the long run. In the setting of this paper, one important adaptation mechanism—timing entry and exit from the fishery—cannot be indefinitely maintained. If climate change permanently pushes fishing grounds so far offshore that entry is no longer profitable, then this adaptation strategy will no longer provide any aid. The question of dynamics in individual adaptation to a changing climate is an important open questions in climate economics.

Looking across firms, these results are encouraging for the prospects of adaptation by other highly mobile industries with ready access to non-climate exposed production processes. The results also inform the potential effectiveness of information as a climate adaptation policy. According to the baseline results, forecast provision has been helpful in mitigating the damage from ENSO in the setting of albacore fishing. It is important to note that rather than indicating that adaptation is "policy-free" in the sense that it will occur without intervention, the results point to the value of policy-driven information provision. Information externalities imply that public provision of forecasts of weather and climate changes can have a positive welfare impacts even if adaptation mechanisms themselves are private.

References

Albacore Working Group (2014). Stock Assessment of Albacore Tuna in the North Pacific Ocean in 2014.

Auffhammer, M. (2018). Quantifying economic damages from climate change. *Journal of Economic Perspectives 32* (4), 33–52.

Auffhammer, M. (2022). Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption. *Journal of Environmental Economics and Management 114*, 102669.

Barnston, A. G., S. Li, S. J. Mason, D. G. DeWitt, L. Goddard, and X. Gong (2010). Verification of the first 11 years of iri's seasonal climate forecasts. *Journal of Applied Meteorology and Climatology 49*(3), 493–520.

Barnston, A. G., M. K. Tippett, M. L. L'Heureux, S. Li, and D. G. DeWitt (2012). Skill of real-time seasonal enso model predictions during 2002–11: Is our capability increasing? *Bulletin of the American Meteorological Society 93*(5), 631–651.

Barreca, A., K. Clay, O. Deschênes, M. Greenstone, and J. S. Shapiro (2016). Adapting to climate change: The remarkable decline in the u.s. temperature-mortality relationship over the 20th century. *Journal of Political Economy 124*(1), 105–159.

Bauer, P., A. Thorpe, and G. Brunet (2015). The quiet revolution of numerical weather prediction. *Nature 525*(7567), 47–55.

Brunner, A. D. (2002). El Niño and world primary commodity prices: Warm water or hot air? *Review of Economics and Statistics 84*(1), 176–183.

Burke, M. and K. Emerick (2016). Adaptation to climate change: Evidence from US agriculture. *American Economic Journal: Economic Policy 8*(3), 106–140.

Cane, M. A., S. E. Zebiak, and S. C. Dolan (1986). Experimental forecasts of El Niño. *Nature 321*, 827–832.

Carleton, T., A. Jina, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, R. E. Kopp, K. E. Mccusker, I. Nath, J. Rising, A. Rode, H. K. Seo, A. Viaene, J. Yuan, and A. T. Zhang (2022). Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *Quarterly Journal of Economics 137*(4), 2037–2105.

Carleton, T. A. and S. M. Hsiang (2016). Social and economic impacts of climate. *Science 353*(6304), aad9837.

Chen, D., M. A. Cane, S. E. Zebiak, R. Canizares, and A. Kaplan (2000). Bias correction of an ocean-atmosphere coupled model. *Geophysical Research Letters 27*(16), 2585–2588.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics 92*(1), 1–45.

Costello, C. J., S. Polasky, and A. Solow (2001). Renewable resource management with environmental prediction. *Canadian Journal of Economics 34*(1), 196–211.

Dell, M., B. F. Jones, and B. A. Olken (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics 4*(3), 66–95.

Dell, M., B. F. Jones, and B. A. Olken (2014). What do we learn from the weather? The new climate-economy literature. *Journal of Economic Literature 52*(3), 740–798.

Deschênes, O. and M. Greenstone (2007). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review 97*(1), 354–385.

Downey, M., N. Lind, and J. G. Shrader (2023). Adjusting to rain before it falls. *Management Science*.

Efron, B. and C. Morris (1975). Data analysis using Stein's estimator and its generalizations. *Journal of the American Statistical Association 70*(350), 311–319.

Enfield, D. B. and A. M. Mestas-Nuñez (2000). Global modes of ENSO and non-ENSO sea surface temperature variability and their associations with climate. In *El Niño and the Southern Oscillation: Multiscale Variability and Global and Regional Impacts*, pp. 89–112. Cambridge University Press.

EPA (2017). Adapting to climate change. 19january2017snapshot.epa.gov/climatechange/adapting-climate-change_.html. Accessed on: 2021/06/11.

Fiedler, P. C. and H. J. Bernard (1987). Tuna aggregation and feeding near fronts observed in satellite imagery. *Continental Shelf Research 7*(8), 871–881.

Glantz, M. H. and M. H. Glantz (2001). *Currents of change: impacts of El Niño and La Niña on climate and society*. Cambridge University Press.

Guo, C. and C. Costello (2013). The value of adaption: Climate change and timberland management. *Journal of Environmental Economics and Management 65*(3), 452–468.

Hamill, T. M. and J. Juras (2006). Measuring forecast skill: is it real skill or is it the varying climatology? *Quarterly Journal of the Royal Meteorological Society 132*, 2905–2923.

Hornbeck, R. and P. Keskin (2014). The historically evolving impact of the ogallala aquifer: Agricultural adaptation to groundwater and drought. *American Economic Journal: Applied Economics 6*(1), 190–219.

Hsiang, S. M., K. C. Meng, and M. A. Cane (2011). Civil conflicts are associated with the global climate. *Nature 476*(7361), 438–41.

Hsiang, S. M. and D. Narita (2012). Adaptation to cyclone risk: Evidence from the global cross-section. *Climate Change Economics 3*(2), 28.

IPCC (2014). Part a: Global and sectoral aspects. (contribution of working group ii to the fifth assessment report of the intergovernmental panel on climate change). *Climate Change 2014: Impacts, Adaptation, and Vulnerability*, 1132.

Kolstad, C. D. and F. C. Moore (2020). Estimating the economic impacts of climate change using weather observations. *Review of Environmental Economics and Policy 14*(1), 1–24.

Kovats, R. S., M. J. Bouma, S. Hajat, E. Worrall, and A. Haines (2003). El niño and health. *Lancet 362*(9394), 1481–9.

Lehodey, P., F. Chai, and J. Hampton (2003). Modelling climate-related variability of tuna populations from a coupled ocean–biogeochemical-populations dynamics model. *Fisheries Oceanography 12*(4-5), 483–494.

Lemoine, D. (2021). Estimating the consequences of climate change from variation in weather. *National Bureau of Economic Research No. 25008*, 38.

Manski, C. F. (2004). Measuring expectations. *Econometrica 72*(5), 1329–1376.

McGowan, J. A., D. R. Cayan, and L. M. Dorman (1998). Climate-ocean variability and ecosystem response in the northeast pacific. *Science 281*(5374), 210–217.

Mendelsohn, R., W. D. Nordhaus, and D. Shaw (1994). The impact of global warming on agriculture: a ricardian analysis. *American Economic Review 84*(4), 753–771.

Meza, F. J., J. W. Hansen, and D. Osgood (2008). Economic value of seasonal climate forecasts for agriculture: review of ex-ante assessments and recommendations for future research. *Journal of Applied Meteorology and Climatology 47*(5), 1269–1286.

Moore, F. C. and D. B. Lobell (2014). Adaptation potential of European agriculture in response to climate change. *Nature Climate Change 4*(7), 610–614.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica 29*(3), 315–335.

NASA (2002). Weather forecasting through the ages. https://earthobservatory.nasa.gov/features/WxForecasting/wx2.php. Accessed on: 2020/3/10.

NCEP (2020). http://www.cpc.ncep.noaa.gov/products/CDB/CDB_Archive_html/CDB_archive.shtml. Accessed on: 2020/3/10.

Neidell, M. M. (2009). Information, avoidance behavior, and health: The effect of ozone on asthma hospitalizations. *The Journal of Human Resources 44*(2), 450–478.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation. *Econometrica 55*(3), 703–708.

NOAA (2013). History of the array. https://tao.ndbc.noaa.gov/proj_overview/taohis_ndbc.shtml. Accessed on: 2021/06/11.

PSMFC (2020). Efin monthly marine fuel prices. www.psmfc.org/efin/data/fuel.html. Accessed on: 2020/3/10.

Roberts, A. (2017). Till tomorrow. https://www.thenewatlantis.com/publications/till-tomorrow.

Rosenzweig, M. and C. R. Udry (2014). Forecasting profitability. *NBER Working Paper 19334*, 47.

Schlenker, W., M. Hanemann, and A. C. Fisher (2005). Will U.S. agriculture really benefit from global warming? accounting for irrigation in the hedonic approach. *American Economic Review 95*(1), 395–406.

Schlenker, W. and M. J. Roberts (2009). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences 106*(37), 15594–8.

Schlenker, W., M. J. Roberts, and D. B. Lobell (2013). U.S. maize adaptability. *Nature Climate Change 3*(8), 690–691.

Schlenker, W. and C. A. Taylor (2021). Market expectations of a warming climate. *Journal of Financial Economics 142*(2), 627–640.

SeaView Fishing (2021). http://www.seaviewfishing.com/News.html. Accessed on: 2021/06/11.

Severen, C., C. Costello, and O. Deschênes (2018). A forward-looking ricardian approach: Do land markets capitalize climate change forecasts? *Journal of Environmental Economics and Management 89*, 235–254.

Sloat, L. L., S. J. Davis, J. S. Gerber, F. C. Moore, D. K. Ray, P. C. West, and N. D. Mueller (2020). Climate adaptation by crop migration. *Nature Communications 11*(1), 1–9.

Stockdale, T. (2021). SEAS5 user guide. Technical report, ECMWF.

Toth, Z. and R. Buizza (2018). *Weather forecasting: What sets the forecast skill horizon?* Elsevier Inc.

Trenberth, K. E. (1997). The definition of El Niño. *Bulletin of the American Meteorological Society 78*, 2771–2777.

Wise, L. (2011). West coast U.S. commercial albacore fishery: Economic analysis.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*, Volume 58. MIT press.

Appendix for online publication

A Model extensions and additional results

A.1 Prices effects

In the model presented in Section 2, expected weather does not affect prices of inputs or outputs. Extending the model to incorporate such effects shows that the definitions and identification results go through with few conceptual changes. Input price changes in particular are inconsequential for the definitions derived in the body of the paper[36] Output price changes potentially confound clean identification of the marginal benefit of adaptation, motivating the empirical analysis in Section A7 and the focus on output rather than revenue effects.

Starting with the same model as in the text, we now allow expected weather to enter both the output price and the input price. The firm knows the forecast and how the forecast affects prices but is still a price taker. To reduce the notation, denote the expected weather as $W_t \coloneqq \mathsf{E}_{t-1}[g(Z_t)]$. The firm solves

$$\max_{x} \mathsf{E}_{t-1}[\pi_t] = p(W_t) f(x_t) W_t - c(W_t) x_t. \tag{13}$$

The first order condition is unchanged in the sense that marginal costs are still equated to marginal benefits of changing each input.

$$p(W_t) f'(x_t) W_t = c(W_t)$$

This implicitly defines the optimal inputs $x_t(p(W_t), c(W_t), W_t)$, and we can find definitions for adaptation, the marginal benefits of adaptation, and the direct effect.

Adaptation is the change in inputs with respect to change in forecasts, which is now

$$A_t \coloneqq \frac{\mathrm{d}x_t(p(W_t), c(W_t), W_t)}{\mathrm{d}W_t} = \frac{\partial x}{\partial p} p'(W_t) + \frac{\partial x}{\partial c} c'(W_t) + \frac{\partial x}{\partial W_t}. \tag{14}$$

The final term is the same as in the main presentation of the model. There are two new terms, both capturing knock-on effects on adaptation that come from price changes. Consider the second term, $\frac{\partial x}{\partial c} c'(W_t)$. Under standard assumptions, $\frac{\partial x}{\partial c} < 0$. If the forecast causes input prices to decline, then $c'(W_t) < 0$, and the whole term will be positive. In other words, if the forecast causes an input to get cheaper, then the firm will engage in more adaptation because the marginal cost of adaptation is

[36] Input price changes in response to forecasts have been documented for agricultural labor markets in Rosenzweig and Udry (2019).

lower.

Taking the derivative of revenue with respect to W_t allows one to define the marginal benefit of adaptation and the direct effect.

$$\frac{\mathrm{d}\mathsf{E}_{t-1}[y_t]}{\mathrm{d}W_t} = \underbrace{p(W_t)f(x_t)}_{\text{Direct eff.}} + \underbrace{p(W_t)f'(x_t)(A_t)W_t}_{\text{M.B. adapt}} + \underbrace{p'(W_t)f(x_t)W_t}_{\text{Output price eff.}}$$

The direct effect is identical to the main text treatment. Condition on forecasts, this effect is identified from the change in ex post revenue with respect to weather realizations. The second pair of terms are what is identified from the effect of forecasts, conditional on realized weather. This involves the marginal benefit of adaptation, including the knock-on effects of price changes on input choices (adaptation). The third term is the output price effect that occurs even if the firm does not change their input choices. To get clean identification of the marginal benefit of adaptation, therefore, it is important to rule out this price effect or to look just at output effects rather than revenue effects.

A.2 Discrete adaptation

The model presented in Section 2 assumed that all adaptation inputs were continuous and that the production function was differentiable in all inputs. These assumptions are not necessary for the formal definition of adaptation, and the estimation strategy presented in the text extends to the case of discrete adaptations. Continuity and differentiability simply help to derive exact expressions for the adaptation decision rule through the implicit function theorem.

In the presence of discrete adaptations, denote adaptation as the vector of changes in inputs with respect to changes in expected weather, or

$$\mathbf{A} = \left(\frac{\Delta x_1(p, \mathbf{r}, \mathsf{E}[g(Z)])}{\Delta \mathsf{E}[g(Z)]}, \ldots, \frac{\Delta x_J(p, \mathbf{r}, \mathsf{E}[g(Z)])}{\Delta \mathsf{E}[g(Z)]} \right)'.$$

In this case, estimation proceeds as in Section 4. For a single input, estimating adaptation can be thought of as estimating the reduced form of an instrumental variables (IV) regression where the first stage is a regression of weather expectations on inputs and the second stage is a regression of inputs on output conditional on realized weather. In this case, the distribution of the input variable is irrelevant to consistent estimation of the reduced form so long as there is identifying variation in weather expectations (Wooldridge, 2010, pg. 84).

This result illustrates, however, that the method presented here cannot be used, in general, to determine the contribution of individual adaptation mechanisms to total adaptation. In an IV setting, one would need as many instruments as inputs to fully identify the effect of each input. Expectations only provide a single instrument that is blunt from the perspective of each individual adaptation mechanism. More importantly, because expectations enter all non-separable inputs, omitting one input from the second stage equation would lead to bias.

Finally, a specific example worth highlighting is the case where a firm has the choice of two possible production functions,

$$y_{it} = \begin{cases} f_1(\mathbf{x}_{it})g(Z) & \text{if } \mathsf{E}[f_1(\mathbf{x}_{it})] \geq \mathsf{E}[f_2(\mathbf{x}_{it})] \\ f_2(\mathbf{x}_{it})g(Z) & \text{if } \mathsf{E}[f_1(\mathbf{x}_{it})] < \mathsf{E}[f_2(\mathbf{x}_{it})] \end{cases}$$

Define the indicator d as $d = \mathbb{1}\{\mathsf{E}[f_1(\mathbf{x}_{it})] \geq \mathsf{E}[f_2(\mathbf{x}_{it})]\}$ and the probability p as $p = P(\mathsf{E}[f_1(\mathbf{x}_{it})] \geq \mathsf{E}[f_2(\mathbf{x}_{it})])$, so output can be written as

$$\begin{aligned} \mathsf{E}[y_{it}] &= \mathsf{E}[df_1(\mathbf{x}_{it})g(Z) + (1 - d)f_2(\mathbf{x}_{it})g(Z)] \\ &= pf_1(\mathbf{x}_{it})\mathsf{E}[g(Z)] + (1 - p)f_2(\mathbf{x}_{it})\mathsf{E}[g(Z)]. \end{aligned}$$

The partial derivative of output with respect to realized weather will be unaffected by this set-up because the weather term can be distributed to the front of the output expression. Moreover, the choice of $\mathbf{x}$ is still a function of $\mathsf{E}[g(Z)]$ in both f_1 and f_2, so the reduced form estimation logic from above applies.

A.3 Mixed input timing decisions

The model presented in Section 2 assumes that all inputs are decided before the random variable Z is realized each period. Here, I relax that assumption.

Consider two inputs, x_1 and x_2, where x_1 is determined before the random variable realizes (which I will call ex ante) and x_2 is determined after the random variable realizes (ex post). Consider a single firm so that entity subscripts can be dropped and normalize the output price to 1. The problem can be solved by backward induction. The firm's ex post problem is

$$\max_{x_{2t}} \pi_t = f(x_{1t}, x_{2t})g(z_t) - p_1 x_{1t} - p_2 x_{2t} \tag{15}$$

given a fixed x_1 from the beginning of the period and a realization, z, of Z. The first order condition is

$$f_2(x_{1t}, x_{2t})g(z_t) = p_2$$

This condition makes clear that x_2 will generally be a function of the realized weather through $g(z)$. In addition, it will be a function of the expected weather through x_1. For instance, in a Cobb-Douglas case with equal factor shares, the firm would like to equalize inputs ex ante, so it would choose x_1 assuming that $g(z) = E[g(Z)]$. ex post, the firm still has incentive to equalize inputs, so it will choose x_2 closer to the ex ante value than in a purely ex post case.

The ex ante value of adaptation given in Equation (4) will be the same, but estimation of this value using realized data will no longer capture all adaptation because

$$\frac{\partial y}{\partial g(z)} = f_2(x_1, x_2)\frac{\partial x_2}{\partial g(z)} + f(x_1, x_2).$$

The second term is the direct effect, as before, but now part of the value of adaptation, $f_2(x_1, x_2)\frac{\partial x_2}{\partial g(z)}$, will be included in the estimate of the direct effect, which will be included in the magnitude of the coefficient on $g(z_t)$. This will serve to attenuate the estimate of the value of adaptation and increase the magnitude of the estimate of the direct effect. Therefore, in a case with both ex ante and ex post adaptation, the effect of forecasts on revenue bounds total adaptation from below, and the effect of realizations conditional on forecasts bounds the direct effect from above.

A.4 Forecast sufficiency under unbiasedness

In Section 4, simple conditions were given for when forecasts will be perfect proxies for private beliefs. Here, I consider alternative assumptions about the information sets of private agents and a public forecaster and derive implications for the use of forecasts as expectation proxies under the assumption of unbiased forecasts. This setting also allows consideration of forecast dynamics.

To simplify the analysis, consider a weather loss function based on the profit maximization problem given in Equation (1). The function describes the profit or output loss that results from realizations of the random variable Z. Denote expected loss as

$$E[L^p(Z_t, \hat{Z}_t, \mathbf{X}(\hat{\mathbf{Z}})_t, \mathbf{p}_t) | G_{t-h}] \tag{16}$$

where we now allow inputs to be a vector and expectations about the future weather are denoted by $\hat{Z}$. $\mathcal{G}_t \in \mathcal{F}$ is the information available to the firm at time t, so this function gives losses due to the h period ahead (or h horizon) forecast. Denote the argument that minimizes Equation (16) in terms of $\hat{Z}_t$ by $s^p_{t|t-h}$, where the superscript p denotes that this is the private firm's value.

Assume that the firm's loss function is symmetric about $Z_t = 0$. Call the loss function a ***Granger loss function*** if either of the two following conditions hold

1. The first derivative of the function, $L^p_1(Z_t, \hat{Z}_t, \mathbf{X}_t, \mathbf{p}_t)$, is strictly monotonically increasing over the range of Z_t and $\bar{f}(Z)$ is symmetric about $Z = s^p$ where $\bar{f}(Z)$ is the conditional distribution of $Z_t - \mathrm{E}[Z_t|\mathcal{G}_{t-h}]$.

2. The distribution of Z, $f(Z)$, is symmetric about $Z = s^p$, is continuous, and is unimodal.

Under either of these conditions, it can be shown that the optimal forecast is $s^p_{t|t-h} = \mathrm{E}[z_t|\mathcal{G}_{t-h}]$ (Granger, 1969). Symmetric loss is limiting but allows for greatly simplified analysis and easier nonparametric identification. The other conditions are more benign. Condition 1 says that there can be no flat regions in the loss function and that the unforecastable component of the stochastic process is elliptical. With positive marginal cost of action or a quadratic loss function, condition 1 will be met. Condition 2 is met by any elliptical distribution.

Now, consider a professional forecaster that minimizes mean squared error (MSE) conditional on the information set $\mathcal{F}_{t-h}$

$$s_{t|t-h} = \underset{\hat{s}}{\operatorname{argmin}}\, \mathrm{E}[(z_t - \hat{s})^2|\mathcal{F}_{t-h}].$$

Solving the minimization problem, one finds that the public forecast in this case is

$$s_{t|t-h} = \mathrm{E}[z_t|\mathcal{F}_{t-h}].$$

Minimization of MSE loss is used in practice by many weather forecasting agencies (Katz and Murphy, 1997).

Patton and Timmermann (2012) show that MSE forecasts have the following properties which will be useful below.

1. Forecasts are unbiased for all h

2. Forecast errors are unpredictable: $\mathrm{Cov}(s_{t+h|t}, x_t) = 0$ for all $x_t \in \mathcal{F}_t$

3. Longer lead forecasts are less precise:

   - $\mathsf{V}(s_{t+h|t}) \quad \mathsf{V}(s_{t+H|t})$ for all $h \quad H$
   - $\mathsf{V}(\varepsilon_{t+h|t}) \quad \mathsf{V}(\varepsilon_{t+H|t})$ for all $h \quad H$ where $\varepsilon_{t+h|t} = z_{t+h} \quad s_{t+h|t}$ is the forecast error

We also need to be able to compare private forecasts to public forecasts. The lemma below says that variance of forecast error is sufficient for comparing forecast quality.

Lemma A.1. *If $G_t \quad F_t$ and $(F_t)_{t \quad 0}$ is strictly monotonic, then there exists a forecast $s_{\quad j t+k}$ such that $\mathsf{V}(\varepsilon_{\quad j t+k}) = \mathsf{V}(\varepsilon^p_{\quad jt})$ for $k \quad 0$.*

Proof. Forecast properties gives us that $\mathsf{V}(\varepsilon_{\quad jt}) \quad \mathsf{V}(\varepsilon^p_{\quad jt}) \quad \mathsf{V}(\varepsilon_{\quad j \quad})$.

Therefore, by continuity there must exist a $k \quad 0$ satisfying the condition. □

Lemma A.2. *For two forecasts $s^1_{t+h|t}$ and $s^2_{t+h|t}$, an agent with a Granger loss function will choose the forecast with lower variance.*

Proof. For condition one, this result holds due to increasing loss for larger deviations in Z. For condition two, the higher variance forecast will create a mean-preserving spread in conditional Z. □

We now provide versions of the forecast sufficiency assumption stated in Section 4. Assume that $G_t \quad F_t$. In other words, that the public forecaster has access to more information than the private firm. Then it is intuitive that the public forecasts are strictly better than the private forecast, and the firm should use the public forecasts.

Proposition A.3. *If the firm loss function or the data generating process satisfies the Granger (1969) conditions and $G_t \quad F_t$, then $s^p_{t+h|t} = s_{t+h|t}$.*

Proof. The Granger conditions imply that $s^p_{t+h|t} = \mathsf{E}[z_{t+h}|G_t]$, so by Lemma A.1 and MSE-forecast property 3, $G_t \quad F_t$ implies

$$\mathsf{V}(\varepsilon^p_{t+h|t}) \quad \mathsf{V}(\varepsilon_{t+h|t})$$

Therefore by lemma A.2, firm loss is minimized by choosing $s^p_{t+h|t} = s_{t+h|t}$. □

Now consider the case where the private firm knows more than the public forecaster: $G_t \ * \ F_t$

To estimate adaptation, we are interested in $\frac{dy}{ds^p}$. If we observed s^p and $G_t \quad F_t$, the chain rule gives

$$\frac{dy}{ds^p} = \frac{\partial y}{\partial s^p} + \frac{\partial y}{\partial s}\frac{\partial s}{\partial s^p}.$$

The question becomes one of how correlated are changes in the two information sets. If the new information enters both G and F, then s and s^p will both change, and the change in the public forecast will again provide good inference for the change in the private forecast. If, however, G grows by gaining information that is already possessed by the private agent, then $\frac{\partial s}{\partial s^p}$ will equal 0.

The last case is when $G_t + F_t$ and $G_t * F_t$. Here, because forecasts based on F_t are public, the firm will incorporate the public forecast into their private information, leading to $\tilde{s}^p_{tj} = g(s^p_{tj}, s_{tj})$. For example, if the agent produces an ensemble forecast by weighting each input forecast by the 1 over its variance (denoted by $w = 1/\sigma^2$), the result would be

$$\tilde{s}^p_{tj} = \frac{(w^p s^p_{tj} + w s_{tj})}{w^p + w}$$
$$\Rightarrow \frac{\partial \tilde{s}^p}{\partial s} = \frac{w}{w^p + w}$$

In general, the more precise the public forecast relative to the private forecast, the closer the researcher would be to capturing the total effect. If the public forecasts are not sufficient for the private beliefs of the agent, the ideal estimation strategy would be to instrument for agent beliefs using the public forecasts.

B Data details and supporting results

B.1 ENSO forecast data

Real-time forecast values are important for identification. I gathered paper records of forecasts issued in by NOAA in the Climate Diagnostics Bulletin (CDB) from June of 1989 until the early 2000s. From the early 2000s to the present, I used and the digital archive maintained by the International Research Institute for Climate and Society (IRI) at Columbia University. The CDB started releasing forecasts in June 1989 and began incorporating the IRI summaries in April 2003.[37] By the year 2000, the number of forecasts incorporated into the Bulletin had grown from 1 to 8.

Figure A1: Example of ENSO forecast issued in the Climate Diagnostics Bulletin

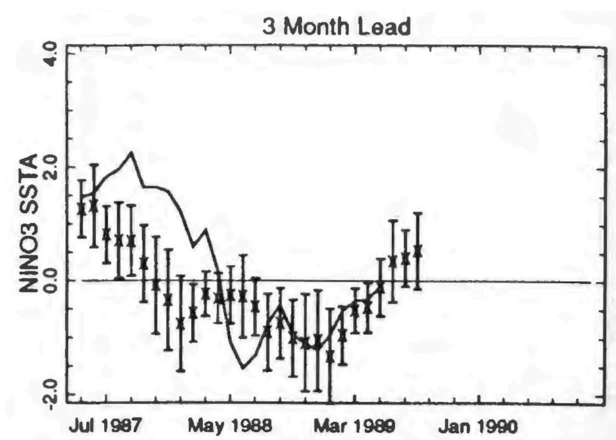

Notes: The figure shows an ENSO forecast issued in the Climate Diagnostics Bulletin in June of 1989. This figure is typical of the forecasts published between 1989 and 2002. The solid line shows the Niño 3 sea surface temperature anomalies and the X are forecasts (and back-casts). Whiskers are the historical standard error for the forecast, a feature present in this but not all models.

To gather the CDB data, I digitized paper records from 1989 to 1999 by scanning each forecast from the Bulletin and then recording the data using the software Graphclick. For Bulletins from 1999 to 2002, I used the online archive of CDBs, again digitizing the figures using Graphclick. For each release, I digitized the Climate

[37]Throughout, I use the 3-month-ahead forecast for estimation. In the June 1989 release of the CDB, three-month ahead forecasts were released, but NOAA also included estimates of the 1 and 2 month-ahead forecasts in the figure (reproduced below as Figure A1). The June 1989 CDB forecasts included data through May 1989, so the Bulletin technically includes a 1-month-ahead forecast for June 1989, a 2-month-ahead forecast for July 1989, and a 3-month-ahead forecast for August 1989. New forecasts in subsequent Bulletins were at the 3-month-ahead horizon during the initial years of publication.

Prediction Center Canonical Correlation forecast (CPC CCA), the Lamont-Doherty Earth Observatory (LDEO) forecasts version 1, 2, and 3; the National Center for Environmental Prediction (NCEP) forecasts, and the Linear Inverse Model (LIM) forecasts. Other forecasts were either issued as maps or contained idiosyncratic issues that prevented digitization.

For data from 2002 through 2016, I used IRI data helpfully supplied to me by Anthony Barnston. These IRI data have formed the basis for analyses of ENSO forecast performance in Barnston et al. (2010, 2012).

In all cases, I used the actual ENSO index values reported in subsequent CDB or IRI reports to calculate forecast accuracy. So, for instance, when digitizing the CPC CCA forecast at a 3 month horizon, I used the actual value reported in the CDB three months later. One could alternatively use a standardized ENSO index across all forecasts. I chose not to do this for a number of reasons. First, all forecasts initially, and many forecasts to the present day, use the Niño 3 index rather than the Niño 3.4 index. Second, the base climatology used to calculate ENSO indices has changed from the 1980s to the present. Third some forecasting agencies might have used their own idiosyncratic calculations of an index or used alternative SST measures. Using the real-time actual values eliminates these sources of noise. On the other hand, what matters for fishing outcomes is the true climate that realized each time period. Thus, for estimation, I use the most recently released version of the Niño 3.4 index. For an alternative method based on scaling alternative index values and visual averaging of maps, see the IRI ENSO *Quick Look*.

B.2 Albacore prices

Albacore prices come from the PacFIN database and are available from 1981 to 2016 at the annual level for ports in the continental United States. Prices are matched to catch using the landing port reported by the vessel.

B.3 Fuel prices

Monthly port-level fuel prices are available for ports in Washington, California, and Oregon from 1999 to the present. The prices are gathered using a phone survey during the first two weeks of each month. The survey respondents are asked to give the price per gallon or price per 600 gallons for number 2 marine diesel before tax.

From 1983 to until the end of 1993, state level prices for number 2 distillate are used for Washington, Alaska, and Oregon. From 1994 until the end of 1998, highway grade number 2 diesel price is used. For Alaska, the state average diesel price is also used for the 1999 to 2016 period.

For California, the distillate price series is not available. State average diesel price is used starting in July of 1995. Prior to July 1995, the gasoline price is used, after accounting for seasonality. In particular, using all data where I observe both gasoline and diesel prices (1994 through 2016) I run the regression

$$\text{diesel}_t = \alpha_{\mathsf{month}} + \gamma_0 \text{gas}_t + \gamma_{\mathsf{month}} \text{gas}_t + \varepsilon_t$$

where *diesel* is the diesel price, *gas* is the gasoline price, α_{month} is a fixed effect for each month of the year $(1, \ldots, 12)$, and $\gamma_{\mathsf{month}} \text{gas}_t$ is an interaction between a fixed effect for each month and the gasoline price. I then predict the diesel price for the pre-1994/5 period using the coefficients from this regression and the observed gasoline price from 1983 to 1995. This procedure should account for intra-year changes in the diesel-gasoline price gap caused by seasonal demand for heating oil. In practice, the seasonal coefficients are not important for this sample.

The same procedure is used to estimate diesel prices for Hawaii over the full sample.

B.4 Vessel movement

Vessel movement is calculated from daily latitude and longitude records plus records of the departure and landing ports. During a fishing trip, movement is calculated as the great circle distance between today's and yesterday's reported location. Calculations were carried out using the `geodist` package in Stata.

For the date of departure, movement is calculated as the great circle distance between the departure port location and the location reported in the first logbook record for the trip. For the final day of the trip, movement is calculated as the great circle distance between the last location reported in the logbook and the landing port.

B.5 Catch weight

Exact catch weight was not recorded in the logbook records for roughly one-third of the daily observations. For the missing records, weight was interpolated in order to obtain complete records for the creation of revenue measures. The interpolation used two methods. First, if a total weight of fish catch was recorded for the trip, then this average weight was used for all fish caught on the trip. For the remaining cases, a regression of weight on gear type, year, and month was used to estimate weight.

Table A1 assesses the effect of this interpolation procedure on the baseline results. Column 1 reproduces the baseline results from Table 2 using only the sub-sample of observations with recorded catch weight. Inference is nearly identical to baseline

Table A1: Robustness to Interpolation of Catch Weight

	(1) Num. fish caught	(2) Catch weight	(3) Catch weight interpolated	(4) Revenue	(5) Num. fish caught
Niño 3.4	-0.072***	-0.074***	-0.056**	-0.11***	-0.049***
	(0.023)	(0.022)	(0.022)	(0.023)	(0.012)
$\widehat{\text{Niño}}$ 3.4	-0.16***	-0.15***	-0.17***	-0.17***	-0.31***
	(0.034)	(0.033)	(0.034)	(0.034)	(0.021)
Covariates	Baseline	Baseline	Baseline	Baseline	Baseline
Weight measure	Observed	Observed	Interpolated	Observed	Observed
Observations	118,692	118,692	120,674	118,692	146,251

Notes: The table shows results from estimating versions of equation (11) on monthly data. The dependent variable in each model is monthly number of fish caught. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects as well as two additional lags of three-month ahead forecasts and realizations of the Niño 3.4 index unless otherwise noted. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987), unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

in this case. Columns 2 and 3 show the baseline regression with catch weight as the dependent variable with and without the interpolation, respectively. One can see that the interpolation increases the magnitude of the results. This occurs because more positive catch observations are being added to the dataset. Column 4 reproduces the revenue result from the baseline table, again showing slightly larger magnitudes but with similar qualitative results between the interpolated and non-interpolated versions. Column 5 shows estimates using the full sample with observed number of fish caught. This is the largest observed sample in the dataset. The effect of forecasts is even stronger in this full dataset than in the sample with observed prices and weight.

B.6 Evidence for Linearity

Figure 1 shows the semiparametric relationship between output and the one-month lag of the Niño 3.4 index in the period before public forecasts existed (from 1981 to June 1989). Both output and the Niño 3.4 index are residualized on baseline controls (year, month-of-year, and vessel fixed effects as well as two additional lags of Niño 3.4). Under the assumption that changes in ENSO relative to the two most recent lags were unforecastable during this period, the plotted relationship recovers the total effect of ENSO on output which, in such a case, would be equal to the direct effect.

From the figure, the relationship between ENSO and output appears to be linear across the range of identifying variation in the Niño 3.4 index. Because this estimate is plausibly unaffected by omitted variable bias from beliefs, it provides evidence for linearity in the direct effect of ENSO on production in this setting.

B.7 Interactions between ENSO and forecasts

The main estimates use a linear specification that does not include interactions between ENSO and ENSO forecasts. Here, I assess the effect of that specification by comparing it to one that includes interactions.

Table A2: Interaction between ENSO and forecasts

	(1) Catch	(2) Catch
Niño 3.4	-0.063***	-0.053**
	(0.024)	(0.025)
$\widehat{\text{Niño}}$ 3.4	-0.19***	-0.19***
	(0.035)	(0.035)
Niño 3.4 $\widehat{\text{Niño}}$ 3.4		-0.014**
		(0.0064)
SEs	Spatial	Spatial
Observations	120,674	120,674

Notes: The table shows results from estimating equation (11) that also includes an interaction between ENSO realizations and forecasts, on monthly data. The dependent, catch, is the total number of fish caught per month. Additional controls are the same as in Table 2. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

The results show that the interaction of the two terms has a significant but practically small effect on catch. The coefficient on the interaction is roughly 1/4 the size of the effect of ENSO realizations and is 1/10 the size of the main forecast effect. Over the range of values consistent with the data, the interaction term has a negligible effect on the conclusions given in the body results section and I therefore maintain the assumption of non-interaction for simplicity or presentation.

B.8 Nonlinear Estimating Equation

Evidence from Figure 1 suggests that a linear specification is reasonable in this setting. *A priori*, however, a nonlinear specification could be reasonable if it is deviations from normal climate in either a hot or cold direction that matter for output. In such a case, a quadratic function for g could approximate the effects of weather.

$$g(z_{t\ 1}) = \gamma_{q;0} + \gamma_{q;1} z_{t\ 1} \quad \gamma_{q;2} z_{t\ 1}^2 \tag{17}$$

With this function of weather, if agents are forming distributional beliefs about ENSO, then the correct forecast term to include would be $g(\hat{z}_{t\ 1}) = \gamma_{q;0} + \gamma_{q;1} \mathsf{E}_{t\ h}[Z_{t\ 1}]$ $\gamma_{q;2} \mathsf{E}_{t\ h}[Z_{t\ 1}^2]$, where h is how far in advance the forecast was issued (at least $h > 1$ in this case). In practice, I observe point forecasts of ENSO, so I will use

$$g(\hat{z}_{t\ 1}) = \gamma_{q;0} + \gamma_{q;1} \mathsf{E}_{t\ h}[Z_{t\ 1}] \quad \gamma_{q;2} \mathsf{E}_{t\ h}[Z_{t\ 1}]^2 \tag{18}$$

This necessitates one of two additional assumptions. Either one can assume that agents are not forming time-varying distributional beliefs about ENSO so that the changes in the point forecast fully capture both linear and nonlinear changes in expectations, or one can assume constant variance of Z. To see the need for the constant variance assumption, assume that agents forecast higher moments of the ENSO distribution. Then

$$\mathsf{E}[g(Z)] = \gamma_{q;0} + \gamma_{q;1} \mathsf{E}_{t\ h}[Z_{t\ 1}] \quad \gamma_{q;2} \mathsf{E}_{t\ h}[Z_{t\ 1}^2] \tag{19}$$

The difference between this value and the measure used for estimation is

$$\mathsf{E}[g(Z)] \quad g(\mathsf{E}[Z]) = \gamma_{q;2}(\mathsf{E}_{t\ h}[Z_{t\ 1}]^2 \quad \mathsf{E}_{t\ h}[Z_{t\ 1}^2]) = \gamma_{q;2} \mathsf{V}_{t\ h}(Z_t) \tag{20}$$

If one assumes that Z_t has constant variance over time, then (20) is constant, and the difference between the two measures will be absorbed by the intercept term. Despite a difference in levels, changes in the two values will carry the same identifying information.

Whether these assumptions limit the interpretation of results is context specific. In C Figure A2, I assess the stability of the variance of ENSO over time. Aside from a period of high variance in the late 1990s, ENSO appears to have a stable second moment. Future work would benefit from using distributional forecasts to assess adaptation to changes in the full distribution of weather.

Figure A2: Second Moments

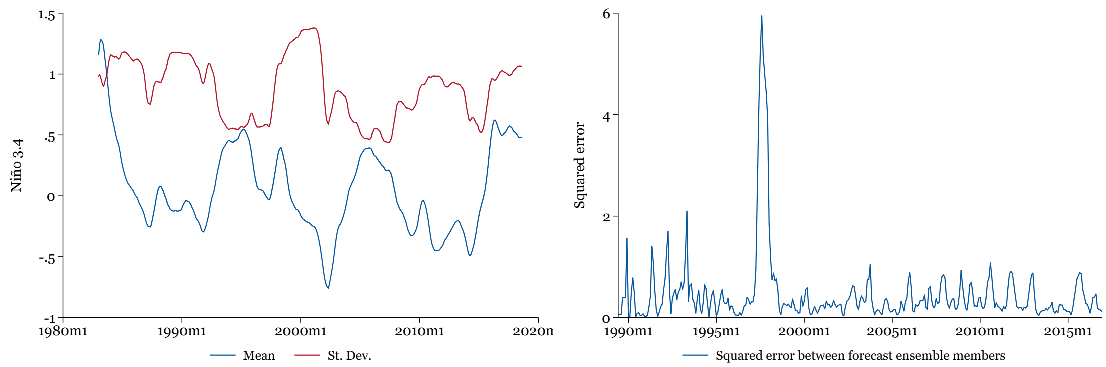

(a) ENSO Rolling St. Dev. and Average (b) Disagreement in Ensemble Members

Notes: Panel (a) shows the moving average and standard deviation of the Niño 3.4 index. Rolling values use a four year window and monthly data. Panel (b) shows the squared error of ensemble members in the ENSO forecast each month.

Putting these elements together, the nonlinear estimating equation is

$$y_{it} = \beta_{q;0} + \beta_{q;1} z_{t-1} + \beta_{q;2} z_{t-1}^2 + \beta_{q;3} \hat{z}_{t-1} + \beta_{q;4} \hat{z}_{t-1}^2 + \mathbf{x}_{it}' \alpha_q + \varepsilon_{q;it} \tag{21}$$

where y_{it} is output or revenue for vessel i at time t, time is measured in months, z_{t-1} is the realized value of the Niño 3.4 index the previous month, $\hat{z}_{t-1}$ is the forecast of ENSO, $\mathbf{x}$ is a vector of control variables (vessel, year, and month fixed effects in the baseline specification), and ε is a stochastic error term. Adaptation is more complicated to assess with this estimating equation and will be considered formally in Section B.9.

B.9 Nonlinear effects of ENSO

Table A3 shows nonlinear effects of ENSO on output and revenue. The left-hand side variable in columns 1 and 2 is output and in columns 3 and 4 it is revenue. Columns 1 and 3 estimate equation (21). Columns 2 and 4 add interactions between the forecast and realization of ENSO. For ease of interpretation, Table A4 shows the marginal effects for each model when both the forecast and realization of ENSO are equal to 1 (moderate El Niño).

The quadratic estimates reinforce the primary results from Table 2. First, in

Table A3: Effect of ENSO on Standardized Output and Revenue: Quadratic Models

		(1) Catch	(2) Catch	(3) Revenue	(4) Revenue
Niño 3.4		-0.070***	-0.082***	-0.11***	-0.12***
		(0.026)	(0.026)	(0.025)	(0.024)
Niño 3.4	Niño 3.4	-0.037***	-0.097***	-0.030***	-0.12***
		(0.011)	(0.027)	(0.011)	(0.026)
Niñò 3.4		-0.19***	-0.19***	-0.17***	-0.18***
		(0.036)	(0.037)	(0.034)	(0.036)
Niñò 3.4	Niñò 3.4	-0.088***	-0.18***	-0.076***	-0.22***
		(0.022)	(0.027)	(0.020)	(0.026)
Niño 3.4	Niñò 3.4		0.16***		0.24***
			(0.054)		(0.053)
Baseline controls		Yes	Yes	Yes	Yes
Unique Vessels		1,214	1,214	1,214	1,214
Observations		120,674	120,674	120,674	120,674

Notes: The table shows results from estimating equation (21) on monthly data. The dependent variable in each model is indicated at the top of the column. All dependent variables are standardized. Catch is the total number of fish caught per month. Revenue is the total ex-vessel value of catch. Additional controls are the same as in Table 2 and are two additional lags of the Niño 3.4 index, two additional lags of forecasts, and fixed effects for vessel, year, and month. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

all quadratic models, the squared terms are significantly different from zero, but the models show that over most of the range of the data, the linear model does a reasonable job capturing the effect of ENSO on the fishery.

Second, both forecasts and realizations of ENSO are important for production in this setting. But conditional on forecasts, realizations are generally an order of magnitude less important than the forecasts themselves. In the context of the model, these estimates indicate that the marginal benefit of adaptation is large compared to the direct effect. The marginal effects show this clearly: the marginal effect of the forecast on output is 6 to 10 times larger than the marginal effect of realized ENSO and 2 to 5 times larger for revenue.

Third, models that do not include forecasts show that as in the linear case, excluding forecasts leads to severe bias.[38] If the forecasts are not included, the direct effect of a moderate El Niño is over-estimated by roughly 100% while the total effect is under-estimated by about 100% as well.

Table A4: Marginal Effects of Quadratic Models at Niño 3.4 and $\widehat{\text{Niño}}$ 3.4 Equal to 1

	(1) Catch	(2) Catch	(3) Revenue	(4) Revenue
Niño 3.4	-0.14***	-0.11***	-0.17***	-0.11***
	(0.037)	(0.035)	(0.034)	(0.032)
$\widehat{\text{Niño}}$ 3.4	-0.37***	-0.39***	-0.33***	-0.38***
	(0.049)	(0.053)	(0.047)	(0.054)
Model	Quadratic	+ Interaction	Quadratic	+ Interaction

Notes: The table shows marginal effects from estimates in Table A3. Standard errors calculated using the delta method. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Finally, the non-zero interaction terms allow for more complex results when forecasts and realizations differ. If the realization of ENSO is 1 and the forecast is unexpectedly high, the direct effect is attenuated and can even turn positive for sufficiently benign conditions. In contrast, if the firm expects conditions to be more benign than an ENSO of 1, then the direct effect is substantially worse. Similar results hold when considering a fixed forecast and changes in realizations of ENSO.

B.10 Evaluating other expectation proxies

If forecasts are not available, other forecast proxies might be elements in the information set of the agent or leads of the right-hand-side variable. The tables below assess the effect of using ***leads*** of ENSO as such proxies. The first table, Table A5, compares the effect of including ENSO forecasts, as in the core results in Section 5.1, versus including a similar-horizon lead. Columns 1 and 2 reproduce baseline estimate results. As shown in the body of the paper, including the forecast reduces the coefficient on the realization of ENSO, and the forecast coefficient itself is large and economically meaningful. Including the 2-month lead, as in Column 3, does reduce the coefficient on the realization of ENSO. If we consider Column 2 to be capturing the “true effect”, then including the lead moves the coefficient on the realization closer

[38] These results are reported in Section B.9.

to the truth. The coefficient on the lead itself is small, so inferring the amount of adaptation from that coefficient will lead—in this case—to a high degree of bias. In theory, the lead might be an unbiased but noisy proxy for agent expectations about future conditions, so we would expect the lead coefficient to be attenuated relative to the true adaptation value coefficient. Finally, column 4 assesses proxy sufficiency. If the forecast is a better proxy of agent beliefs than the lead, it should "out compete" the lead when it comes to explaining firm output (see Section 2 for more details). Column 4 shows that this is indeed the case. When including both the forecast and lead, the coefficient on the lead drops substantially, while the forecast coefficient is not affected appreciably.

Table A6 assesses the effect of including leads of different horizons. The first column includes the 1-month-ahead lead, and each column moves the lead one more month into the future. Column 2 is the same as Column 3 from Table A5. One can see that as the lead moves further into the future, the coefficient on the lead gets smaller and smaller. The lead is acting as a progressively worse proxy for agent beliefs, again under the assumption that the baseline results that include forecasts are accurate. Concomitantly, the coefficient on the realization of ENSO gets progressively larger and larger. With the three and four-month-ahead leads, the coefficient on the realization is back up to roughly the same level as one observes when not including any agent belief proxy.

Table A5: Comparing Effect of ENSO Forecasts and ENSO Leads

	(1) Catch	(2) Catch	(3) Catch	(4) Catch
Niño $3.4_{t\ 1}$	-0.091***	-0.063***	-0.070***	-0.055**
	(0.022)	(0.024)	(0.026)	(0.027)
Niño $3.4_{t\ 1}$		-0.19***		-0.19***
		(0.035)		(0.034)
Niño 3.4_{t+2}			-0.029*	-0.011
			(0.016)	(0.015)
Baseline controls	Yes	Yes	Yes	Yes
Unique Vessels	1,214	1,214	1,214	1,214
Observations	120,674	120,674	120,674	120,674

Notes: The table shows results from estimating versions of equation (12) on monthly data. The dependent variable in each model is total catch in the month. All dependent variables are standardized. Additional controls are the same as in Table 2 and are two additional lags of the Niño 3.4 index, two additional lags of forecasts (Columns 2 and 4), and fixed effects for vessel, year, and month. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Table A6: Illustrating Attenuation When Using ENSO Leads

	(1) Catch	(2) Catch	(3) Catch	(4) Catch
Niño 3.4$_{t \quad 1}$	-0.041	-0.070***	-0.086***	-0.091***
	(0.029)	(0.026)	(0.025)	(0.022)
Niño 3.4$_{t+1}$	-0.052***			
	(0.017)			
Niño 3.4$_{t+2}$		-0.029*		
		(0.016)		
Niño 3.4$_{t+3}$			-0.013	
			(0.020)	
Niño 3.4$_{t+4}$				-0.0042
				(0.023)
Baseline controls	Yes	Yes	Yes	Yes
Unique Vessels	1,214	1,214	1,214	1,214
Observations	120,674	120,674	120,674	120,674

Notes: The table shows results from estimating versions of equation (12) on monthly data. The dependent variable in each model is total catch in the month. All dependent variables are standardized. Additional controls are the same as in Table 2 and are two additional lags of the Niño 3.4 index and fixed effects for vessel, year, and month. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 24 months of lags for autocorrelation (Conley, 1999, Newey and West, 1987). Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

C Additional figures and tables

Figure A3: ENSO Cycle

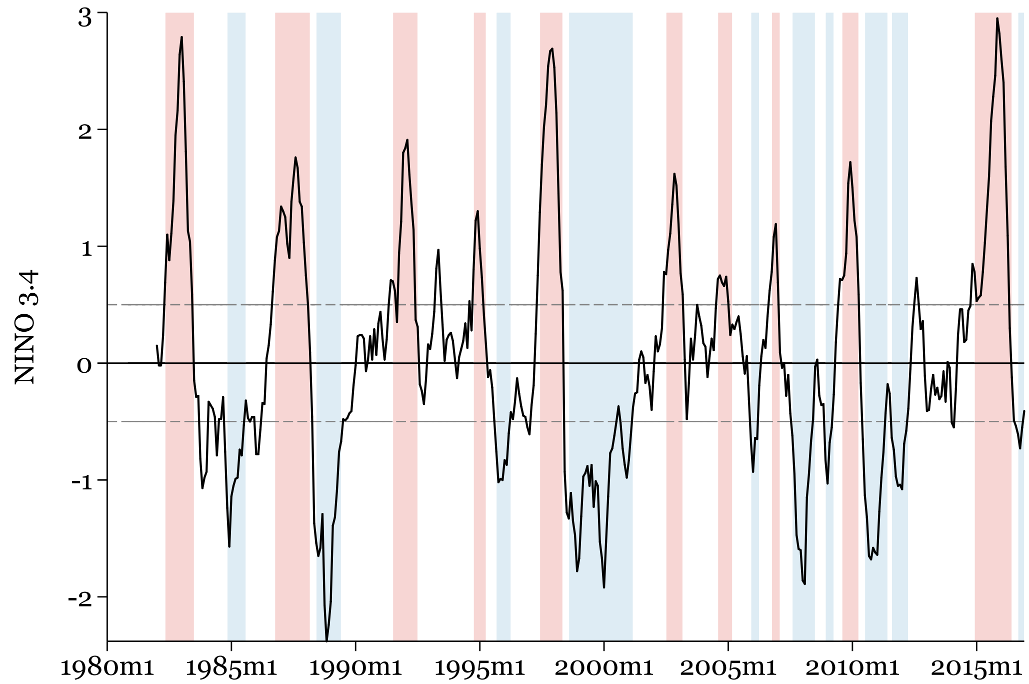

Notes: The ENSO cycle is measured here by the Niño 3.4 index, which is the three month moving average of sea surface temperature anomalies from the Niño 3.4 region of the equatorial Pacific Ocean. Values above 0.5 indicate an El Niño and values below -0.5 indicate La Niña, as denoted by the red and blue shaded regions respectively. For more information on ENSO, see Section 3.

Figure A4: ENSO Forecast Skill

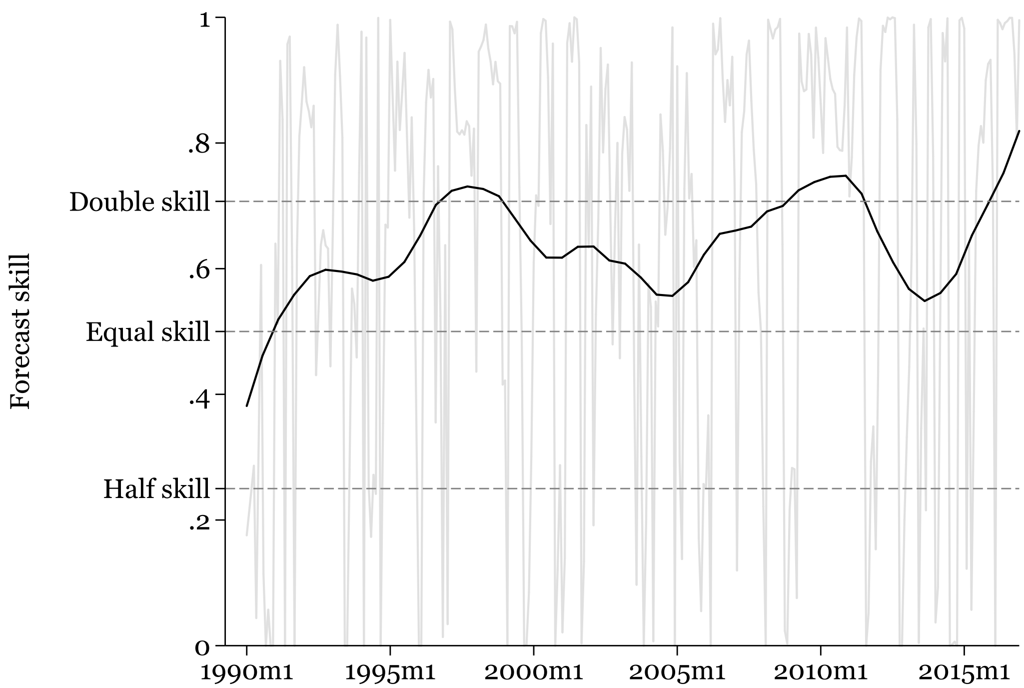

Notes: Forecast skill measured by a normalized version of the Brier skill score is indicated by the light gray line. Skill is the exponential of log of 0.5 times squared error of the forecast divided by the squared error of a naïve persistence forecast. The moving average of monthly skill is given by the black line. The moving average is calculated using a local polynomial regression (Epanechnikov kernel with bandwidth of 12 months). The gray, dashed lines indicate different levels of forecast quality. The bottom line is where the professional forecast has twice as high of standard error as a persistence forecast. The middle line is where the two forecasts are of equal quality. The top line is where the professional forecast has half the standard error of the persistence forecast.

Figure A5: Correlation Between Niño 3.4 and Sea Surface Temperature

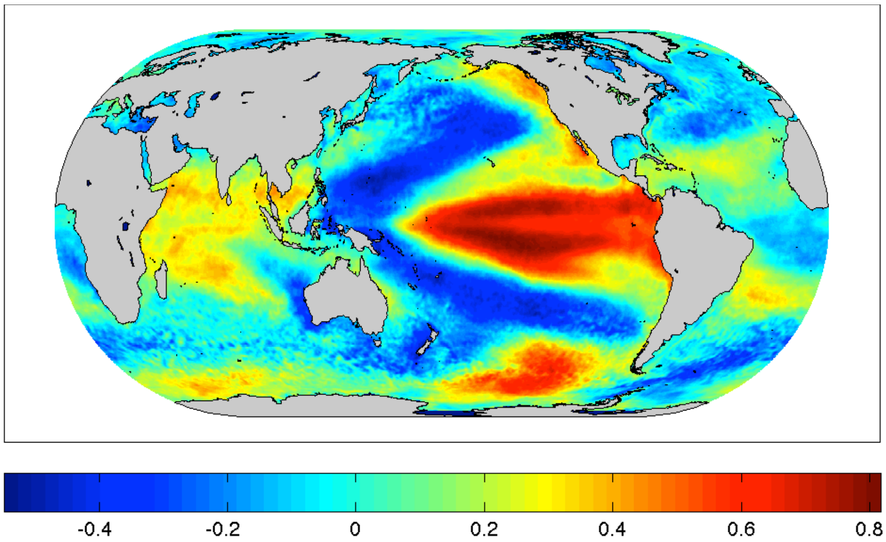

Notes: The heat map shows correlation between the one month lag of the Niño 3.4 index and sea surface temperature for each quarter degree latitude-longitude grid cell.

Figure A6: Fishing and Transiting Locations for Daily Observations

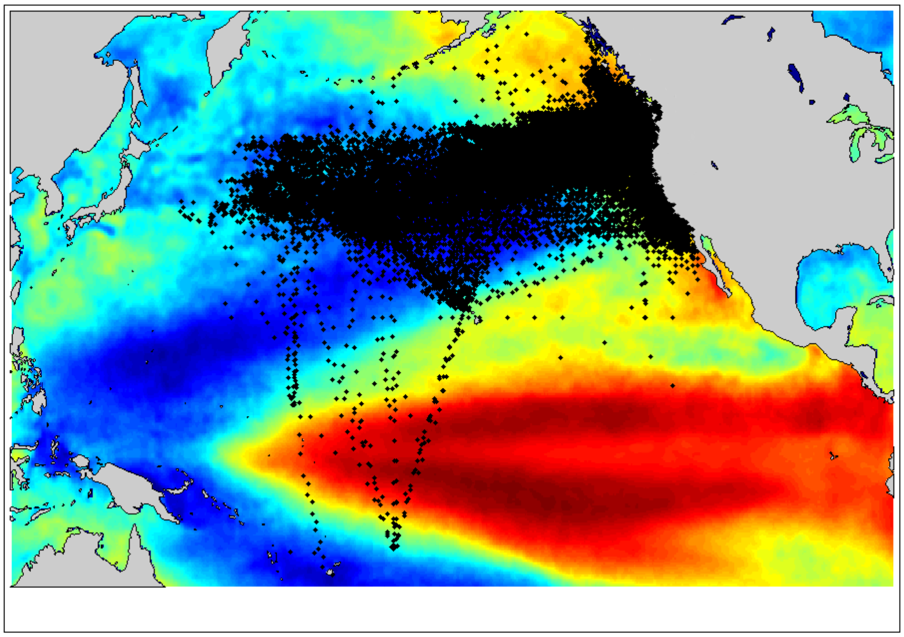

Notes: The heat map shows correlation between the one month lag of the Niño 3.4 index and sea surface temperature for each quarter degree latitude-longitude grid cell, as in Figure A5. Each point shows a daily observation of either fishing or transiting for a subset of the data from 1981 to 2010.

Table A7: Association of ENSO and Albacore Prices

	ln(albacore price)
Niño 3.4 (t-1)	0.015
	(0.035)
Niño 3.4 (t-1)	-0.037
	(0.037)
L.ln(albacore price)	1.03***
	(0.12)
Observations	29

Notes: The table shows results from estimating Newey-West regressions on annual time series data. Given the annual reporting of albacore prices, the table reports the closest time series analogue to the main estimation equation (11). The dependent variable is the log of the wholesale albacore price. In parentheses are Newey-West standard errors with 2 (annual) lags for autocorrelation. Significance indicated by: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A8: Association of ENSO and Input Prices

	(1) ln(diesel price)	(2) ln(wage)
Niño 3.4 (t-1)	0.0031	0.047
	(0.0070)	(0.035)
Niño 3.4 (t-1)	-0.0023	-0.069**
	(0.0096)	(0.033)
Observations	267	107
Frequency	Month	Quarter

Notes: The table shows results from estimating Newey-West regressions on time series data. The model is the closest time series analogue of the main estimation equation (11). The dependent variable in Column 1 is the log of the monthly average fuel price (marine diesel). In Column 2, it is the log of average weekly wage for NAICS 1141 in counties in California, Oregon, and Washington from the QCEW. Both models include 1 lag of the dependent variable. In parentheses are Newey-West standard errors with 24 monthly lags (Column 1) or 8 quarterly lags (Column 2) for autocorrelation. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Table A9: Robustness Set 1 for Quadratic Specification: Marginal Effects

	(1) Vessel by year FEs	(2) Vessel by month FEs	(3) Vessel trends	(4) Nino 3.4 t 12	(5) 6 lags Nino 3.4
Niño 3.4	-0.15***	-0.15***	-0.15***	-0.15***	-0.18***
	(0.033)	(0.029)	(0.037)	(0.038)	(0.040)
Niño^2 3.4	-0.38***	-0.36***	-0.38***	-0.40***	-0.30***
	(0.042)	(0.040)	(0.049)	(0.058)	(0.051)
Observations	120,301	118,919	120,674	112,908	118,982

Notes: The table shows marginal effects, evaluated at Niño 3.4 and the forecast of Niño 3.4 equal to 1, from estimating the quadratic version of equation (11) on monthly data. The dependent variable in each model is monthly number of fish caught. All additional covariates, standard errors, and sample are the same as the baseline specification unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Table A10: Robustness Set 2 for Quadratic Specification: Marginal Effects

	(1) Year-month clustering	(2) Full catch sample	(3) Less than 46	(4) Drop 1997 to 2001	(5) Catch lag covariate
Niño 3.4	-0.14*	-0.099**	-0.13***	-0.14***	-0.10***
	(0.075)	(0.046)	(0.036)	(0.040)	(0.030)
Niño^2 3.4	-0.37***	-0.46***	-0.35***	-0.39***	-0.42***
	(0.11)	(0.062)	(0.049)	(0.051)	(0.048)
Observations	120,674	146,251	118,923	91,527	120,674

Notes: The table shows marginal effects, evaluated at Niño 3.4 and the forecast of Niño 3.4 equal to 1, from estimating the quadratic version of equation (11) on monthly data. The dependent variable in each model is monthly number of fish caught. All additional covariates, standard errors, and sample are the same as the baseline specification unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Table A11: Robustness Set 1 for Quadratic Interaction Specification: Marginal Effects

	(1) Vessel by year FEs	(2) Vessel by month FEs	(3) Vessel trends	(4) Nino 3.4 t 12	(5) 6 lags Nino 3.4
Niño 3.4	-0.12***	-0.12***	-0.12***	-0.11***	-0.15***
	(0.031)	(0.028)	(0.035)	(0.036)	(0.037)
$\widehat{\text{Niño}}$ 3.4	-0.40***	-0.37***	-0.41***	-0.44***	-0.31***
	(0.046)	(0.044)	(0.053)	(0.066)	(0.055)
Observations	120,301	118,919	120,674	112,908	118,982

Notes: The table shows marginal effects, evaluated at Niño 3.4 and the forecast of Niño 3.4 equal to 1, from estimating the quadratic interaction version of equation (11) on monthly data. The dependent variable in each model is monthly number of fish caught. All additional covariates, standard errors, and sample are the same as the baseline specification unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Table A12: Robustness Set 2 for Quadratic Interaction Specification: Marginal Effects

	(1) Year-month clustering	(2) Full catch sample	(3) Less than 46	(4) Drop 1997 to 2001	(5) Catch lag covariate
Niño 3.4	-0.11*	-0.068	-0.097***	-0.11***	-0.10***
	(0.067)	(0.043)	(0.034)	(0.036)	(0.030)
$\widehat{\text{Niño}}$ 3.4	-0.39***	-0.48***	-0.36***	-0.40***	-0.36***
	(0.12)	(0.055)	(0.053)	(0.055)	(0.048)
Observations	120,674	146,251	118,923	91,527	120,674

Notes: The table shows marginal effects, evaluated at Niño 3.4 and the forecast of Niño 3.4 equal to 1, from estimating the quadratic interaction version of equation (11) on monthly data. The dependent variable in each model is monthly number of fish caught. All additional covariates, standard errors, and sample are the same as the baseline specification unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Appendix References

Barnston, A. G., S. Li, S. J. Mason, D. G. DeWitt, L. Goddard, and X. Gong (2010). Verification of the first 11 years of iri's seasonal climate forecasts. *Journal of Applied Meteorology and Climatology 49*(3), 493–520.

Barnston, A. G., M. K. Tippett, M. L. L'Heureux, S. Li, and D. G. DeWitt (2012). Skill of real-time seasonal enso model predictions during 2002–11: Is our capability increasing? *Bulletin of the American Meteorological Society 93*(5), 631–651.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics 92*(1), 1–45.

Granger, C. W. J. (1969). Prediction with a generalized cost of error function. *Operations Research Quarterly 20*(2), 199–207.

Katz, R. W. and A. H. Murphy (1997). *Economic value of weather and climate forecasts*. Cambridge University Press.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation. *Econometrica 55*(3), 703–708.

Patton, A. J. and A. Timmermann (2012). Forecast rationality tests based on multi-horizon bounds. *Journal of Business & Economic Statistics 30*(1), 1–17.

Rosenzweig, M. R. and C. R. Udry (2019). Assessing the benefits of long-run weather forecasting for the rural poor: Farmer investments and worker migration in a dynamic equilibrium model. Technical report, National Bureau of Economic Research.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*, Volume 58. MIT press.