Improving Climate Damage Estimates by Accounting for Adaptation

Jeffrey Shrader*

June 17, 2021

Abstract

Climate change is projected to severely damage the global economy. Adaptation in response to a changing climate will affect how much damage ultimately occurs. A major source of uncertainty in existing damage estimates is the extent to which they include or exclude such adaptation. This paper shows how to estimate damages while accounting for forward-looking adaptation using an empirical strategy that compares responses to forecasts and to realizations of weather. The empirical strategy is applied to study climate damages and adaptation in a fishery using a novel dataset of climate forecasts and detailed, firm-level records. Empirically, the benefit of forward-looking adaptation is large and important, and accounting for adaptation substantially changes the estimated damages from climate variation. Mechanism analysis shows that firms adapt by reducing production costs and timing entry into the fishery. (JEL:D22,D84,Q22,Q54)

*Columbia University, 420 West 118th Street, New York NY 10027 (e-mail: jgs2103@columbia.edu). Special thanks go to Debra Losey from the NOAA Fisheries library for helping me locate historical ENSO forecasts, to Tony 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, 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? A popular approach to answering this question uses panel data and high frequency variation to estimate impacts of weather on economic outcomes, then combines the estimates with climate change projections to predict future damages.[1] One of the primary obstacles to this approach is adaptation—actions taken by individuals to prepare for or adjust to a changing climate. The concern is that using high-frequency variation excludes adaptation, leading to overestimates of the damage from persistent climate change.[2] On the other hand, some argue that this approach returns causal estimates of the direct effect of climate purged of potential confounders including endogenous responses and thus provides a good starting point for evaluating climate damages.[3] In this paper, I show that for typical empirical applications, both arguments are incomplete. The use of forecasts leads agents to engage in forward-looking adaptation that still generally biases estimates based on even the highest frequency data. Accounting for expectations when estimating weather impacts can both reduce the bias in direct effect estimates and allows for estimation of the benefit of adaptation.

The long time horizons involved with climate change creates many challenges when studying the issue. Estimating damages is no exception. The effects of acute shocks can be estimated using short-run variation, but such variation sheds less light on responses to more durable changes in climate. Accurate estimates of external damages from climate change are crucial, however, for optimal policy design. One potential way to make progress on the problem is to combine causal estimates of direct effects of weather with estimates of the costs and benefits of adaptation.[4]

Both of these quantities are challenging to estimate. Expectations bias direct weather effect estimates because people can act on their expectations to minimize risks. Such bias is particularly likely to occur in the context of weather due to the wide availability of high-quality forecasts.[5] Existing methods for estimating weather effects aim to eliminate adaptation by using high-frequency variation. But if people

[1]See Dell et al. (2014) and Auffhammer (2018b) for reviews.

[2]Burke et al. (2016) and Carleton and Hsiang (2016) in reviewing evidence on economic damages from climate change both list better empirical estimates of adaptation as key missing elements. Mendelsohn et al. (1994) made this argument in an important early contribution to the empirical study of climate damages.

[3]A version of this argument based on the envelope theorem is formalized in Hsiang (2016). Deschênes and Greenstone (2007) provided an important early empirical analysis along these lines.

[4]Such estimates could be used to inform integrated assessment models (IAMs) that account for adaptation or a modern climate-macro model that includes adaptation like Fried (2021).

[5]As well as features of the weather system like substantial autocorrelation in many weather processes.

are routinely forming and acting on expectations of current weather conditions, then adaptation will still potentially affect estimation. Envelope theorem-based methods, formalized in Hsiang (2016), seek to purge all adaptation from direct effect estimates by examining dependent variables that are value functions (maximized profit, utility, etc.) and exploiting the equivalence between marginal adaptation benefits and costs at optimum. But this equivalence relies on all adaptation mechanisms being continuous, and many relevant outcomes in climate economics are not value functions.[6] Luckily, a solution to expectation-based bias is conceptually straightforward: include a measure of forecasts used by agents in the estimating equation. The forecasts condition forward-looking adaptation out of the equation and therefore remove it as a source of bias from weather effect estimates.

Including forecasts in the estimating equation has the added advantage of allowing one to estimate the benefits of forward-looking adaptation. Estimating the overall benefits of adaptation is typically hard because many individual mechanisms—choosing different inputs, altering consumption patterns, switching technology—might help reduce damage from a changing environment.[7] For policy, we would like to know the damage that results from changes in the environment after *all* relevant adaptation mechanisms have been incorporated as well as the aggregate cost of those adaptations. Thus, identification of the overall benefit of adaptation either requires *a priori* knowledge of each adaptation mechanism available to agents and suitable exogenous variation for each one, or it involves finding a way to estimate the overall effect of adaptation without reference to the underlying mechanisms.[8] Including forecasts in the estimating equation takes the latter route: using forecasts, the marginal benefit of all *ex ante* adaptation engaged in by the agent can be identified.

Intuitively, identification comes from exploiting the difference between news about future fluctuations versus the realization of those fluctuations. Consider the case of firm production. News, measured by forecasts conditional on realizations of weather, causes the firm to alter inputs in preparation for upcoming fluctuations. Changes in

[6] Also, as Fisher et al. (2012) illustrates for the case of agricultural storage and Lemoine (2021) shows for a wide range of cases, dynamics can still lead to adaptation-based bias when all of the conditions for the static envelope theorem are met.

[7] An extensive literature has shown that individuals and firms employ a number of different adaptation mechanisms to respond to environmental changes. For some of the many examples, see Graff Zivin and Neidell (2012), Barreca et al. (2016), Deschênes et al. (2017), Taraz (2017), Sloat et al. (2020), and Colmer (2021).

[8] Envelope theorem-based methods are also unable to help here because they isolate the direct effect of weather by using the fact that adaptation benefits and costs are equal on the margin to wash out adaptation effects. This necessarily means that the marginal benefits and costs of adaptation cannot be estimated.

inputs are adaptation. The effect of news on revenue is equivalent to the effect of news on profit, so the marginal benefit of all *ex ante* adaptation can be estimated by looking at how firm revenue is affected by news about an upcoming environmental change. And, as described above, the realizations conditional on forecasts identify the effect of weather purged of *ex ante* adaptation.

Estimation using both forecasts and realizations of weather has a few unique benefits when it comes to understanding climate damages. First, as described above, the strategy allows one to more accurately estimate weather effects be removing confounding from forward-looking adaptation. Second, forward-looking adaptation in particular is identified. Adaptation that occurs in advance of a change in the environment could be especially valuable in many environmental contexts—including climate change—where disaster might result from a failure to take precautionary measures. Previous studies of climate impacts have almost uniformly ignored the benefits of forward-looking adaptation.[9] Third, the marginal benefit estimates generated using this strategy can be used to bound marginal costs of adaptation from above, using a generalization of the envelope theorem argument to cases with potentially discontinuous adaptation mechanisms.

Perhaps most importantly, 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.[10] This class of methods has two limitations. First, the results from this paper show that adaptation can still affect estimates based on

[9] Severen et al. (2018) is a notable exception that brings expectations into the Ricardian framework. Despite the abundance of evidence showing that people engage in individual adaptation mechanisms, evidence on overall benefits of adaptation is mixed. See Auffhammer (2018b) for a recent review. One reason could be that existing methods typically use low frequency variation in realized weather to estimate adaptation and thus might miss forward-looking adaptation that occurs before weather is realized. More fundamentally, some researchers have even questioned whether individuals will engage in any *ex ante* adaptation in real-world settings (Mendelsohn, 2000, Repetto, 2008, Massetti and Mendelsohn, 2018). The empirical strategy presented here allows for quantification of the degree of forward-looking adaptation, and the empirical results show that such adaptation is practically important.

[10] 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 et al. (2009, 2012), Hsiang and Narita (2012), Butler and Huybers (2013), Schlenker et al. (2013), Moore and Lobell (2014), Burke and Emerick (2016), Bento et al. (2020). A related method compares estimates derived from high-frequency variation across entities or characteristics as in Auffhammer (2018a) and Carleton et al. (2020).

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.

I apply the empirical strategy presented above to estimate the effects of climate variation from El Niño/Southern Oscillation (ENSO) on albacore tuna harvesters in the North Pacific. The empirical setting is particularly well suitable to studying forecasts and climate damage. ENSO is a major source of global climate variation stemming from periodic but stochastic warming and cooling of the equatorial Pacific Ocean. ENSO was believed to be unforecastable as recently as the mid-1980s. Within the decade, however, breakthroughs in modeling, computing, and data collection allowed climatologists to create accurate forecasts of ENSO multiple months in advance. Concurrent with these developments, the National Oceanic and Atmospheric Administration (NOAA) began a program to disseminate ENSO forecasts to fisheries. The albacore fishery, historically a setting where output and profit declined substantially during ENSO events, was one such fishery. Because the fishery is spatially distant from the area where ENSO forms, these forecasts and attendant NOAA reports on ocean conditions are plausibly the main source of ENSO information available to albacore harvesters over the sample period.

Estimates show that the information in ENSO forecasts is important to the fishery. Forecasts have a much stronger effect on output and revenue than do realizations of ENSO. The effect of an ENSO forecast on output, for example, is three times larger than the effect of realizations. Interpreting this through the lens of the model, the estimates suggest that the benefit of forward-looking adaptation is large relative to the direct effect of ENSO that occurs conditional on that adaptation.

The results also show that if adaptation were ignored when estimating the effect of ENSO (by excluding forecasts from the regression), then estimates would be biased in two ways. First, the direct effect of ENSO would be biased because positive correlation between forecasts and realizations causes some of the adaptation effect to be attributed to the direct effect. Second, the total effect of ENSO would be understated because realizations of ENSO do not capture the full cost of forward-looking adaptation that is identified by variation in expectations. Both sources of bias can lead to improperly set policy if the estimates were to be used to inform a Piguovian tax or other regulation.

Exploiting the richness of the spatially explicit, high-frequency, firm-level data, secondary results examine the mechanisms the vessels use to adapt. Overall, vessels respond to the forecasts by reducing their fishing effort and expenditures during adverse periods. On the intensive margin, in anticipation of ENSO, harvesters fish fewer hours per days, move less during fishing trips, and employ fewer fishing lines. 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 (2014a) 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. Together with the headline estimates, these results highlight both the opportunity and limitations of using information as a public policy response to environmental changes.

Overall, the results show that information has enabled substantial adaptation to climate variation from ENSO in the north Pacific albacore fishery. The same method is potentially widely applicable because ENSO is a major source of climate variation around the world.[11] And ENSO realizations have been shown to affect many important economic outcomes including intensification of global conflict (Hsiang et al., 2011), changes in global commodity prices (Brunner, 2002), changes in agricultural productivity (Rosenzweig and Udry, 2014b), 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.

More broadly, expectations of weather likely affect many existing estimates across multiple fields in economics. Routinely updated, short-range weather forecasts have

[11] Glantz et al. (2001) calls ENSO "the second most important climate process after the changing seasons" due to its widespread effects on the global climate and well-documented impacts on economic and ecological processes.

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 (Barnston et al., 2010, Toth and Buizza, 2018). Even without modern forecasting technology, people have been forming expectations of weather for millennia (NASA, 2002).[12] 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 lag and leads of weather can all be useful ways to reduce expectation-based confounding.[13] Including a forecast is useful when either: (1) expectations are time varying or (2) one wants to explicitly estimate the marginal benefit of adaptation rather than simply treating adaptation as a nuisance parameter.

Finally, forecasts of environmental processes offer a rich and useful context for studying the role of information in economic decisions, an area that is of growing interest in economics (Kamenica, 2017, Haaland et al., 2021). Forecasts of environmental processes are well suited to study this issue not only because they are routinely used by individuals and are readily observable by the researcher, but also because the variation in the underlying environmental processes is often economically exogenous.[14] This feature contrasts with other settings, like finance, where forecasts have the potential to endogenously change the state, complicating empirical analyses. In the future, growing bodies of data and falling costs of data analysis imply that more firms will be making expectation-driven investments, increasing the need and opportunity to study forward-looking behavior.

The rest of the paper proceeds as follows: Section 2 formalizes the role of expecta-

[12]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."

[13]From the perspective in this paper, widely used individual and seasonal fixed effects can be viewed as one method for reducing expectation-based confounding. If individuals form expectations based on long-run averages, then such fixed effects will remove much of the confounding from *ex ante* adaptation. Fixed effects are not the only possible proxy. Lemoine (2021) uses leads of weather as a forecast proxy (estimates using leads versus ENSO forecasts in this paper's context are shown and discussed in Section B.10). Downey et al. (2021) uses the relationship between ENSO and rainfall in the U.S. to estimate monthly rainfall forecasts.

[14]Although, as this paper shows, expectations must be accounted for to arrive at econometric exogeneity.

tions in adaptation, provides conditions under which public forecasts can act as good proxies for agent expectations, and shows that a regression framework can identify both climate adaptation and direct weather effects. 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 and tests of assumptions. 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. Finally, Section 8 concludes.

2 Identifying adaptation and damages

2.1 Expectations identify ex ante adaptation

Economic adaptation is commonly defined as the actions taken by an individual or group of individuals to help reduce the negative effects of a potential change in the environment or to capitalize on gains from such a change.[15] Formalizing this notion of adaptation in a simple production model helps clarify how to estimate both adaptation and the direct impact from environmental shocks. In particular, a formal definition of adaptation will generalize from the single adaptation strategies or mechanisms—staying indoors on hot or polluted days (Neidell, 2009), changing the mix of crops or the use of agricultural inputs (Rosenzweig and Udry, 2014a, Hornbeck and Keskin, 2014), or use of air conditioning (Barreca et al., 2016)—to the overall effect of adaptation on an individual's welfare.

Consider a firm producing a univariate output at time t which is a function of random weather as well as inputs that are chosen by the firm. Assume that the firm's choices only affect profit this period.[16] To emphasize the uncertain effect of weather on the production process, assume that the firm must choose each input, x_{jt}, before the weather in period t is realized.[17] At the beginning of each period, the firm's

[15]For examples of such a definition, see EPA (2017) or IPCC (2014). This study will focus on adaptation by a single economic agent.

[16]In the empirical results, outcomes are analyzed at the monthly level, and harvesters rarely take trips lasting for more than a month. 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. In other words, the model assumes that the problem faced by the firm resets each period, an assumption that appears reasonable in many fisheries (Costello et al., 2001).

[17]For the more general model considering inputs chosen after weather has realized, see Section A. At the end of this section, I discuss which identification results hold for the entirely *ex ante* model presented here versus the more general model. For an extension to the case with finite adjustment costs, see Downey et al. (2021).

problem is to maximize expected profit

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

Output prices are denoted by p, $\mathbf{c}$ is the J-dimensional vector of input prices, $\mathbf{x}$ is the J-dimensional vector of inputs, and Z is a stochastic weather variable with at least one finite moment.[18] Further assume that $f(\mathbf{x})$ is twice continuously differentiable and concave.[19] As is standard, a subscript on an expectation operator denotes the information set on which the expectation is conditioned, so $\mathbb{E}_{t-1}[g(Z_t)] := \mathbb{E}[g(Z_t)|\mathscr{F}_{t-1}]$ is the expected weather this period conditional on information about the weather in all time periods up to and including the most recent period.

Denote realized revenue by $y_t = pf(\mathbf{x}_t)g(z_t)$ and *ex ante* revenue as the expectation of this term with respect to information at $t-1$. Prices are assumed to be constant. In a more general discussion of climate change impacts, it might be appropriate to consider prices that are a function of the climate. The estimator of total adaptation used here will be unaffected by allowing for climate-driven output price changes under additional assumptions on the elasticity of demand for the firm's output that would rule out extra risk taking during adverse events (Allen et al., 2016).[20]

An optimizing firm chooses inputs to maximize the value of Equation (1). Aside from the weather variable, the problem is a standard one, as indicated by the representative first order condition.

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

Adaptation, as per the above definition, is the response of agents to anticipated changes in environmental conditions. In the context of the model, the agent chooses inputs, and environmental conditions are determined by the distribution of weather.

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

[18] For an extension of the model to production functions of the form $f(x, z)$, see Section A. Multiplicative separability is assumed here for ease of presentation. Also, the model is presented with a single weather variable, Z, but nothing prevents the inclusion of a vector of weather variables. In that case, the vectors of derivatives given below would be replaced by Jacobian matrices. I use a vector of inputs to emphasize the potentially high dimensionality of the firm's adaptation choices (and the attendant econometric challenge).

[19] 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.

[20] In the empirical setting, the assumption of prices being uncorrelated with weather is testable and appears to hold. See Section C.

implicitly defined by Equation (2) can be denoted $x^*_{jt}(p, \mathbf{c}, \mathbb{E}_{t-1}[g(Z_t)])$ for all j and t, so the formal definition of adaptation is

$$\mathbf{A}_t = \left(\frac{\partial x^*_{1t}(p, \mathbf{c}, \mathbb{E}_{t-1}[g(Z_t)])}{\partial \mathbb{E}_{t-1}[g(Z_t)]}, \ldots, \frac{\partial x^*_{Jt}(p, \mathbf{c}, \mathbb{E}_{t-1}[g(Z_t)])}{\partial \mathbb{E}_{t-1}[g(Z_t)]}\right)' = \frac{\partial \mathbf{x}^*_{\mathbf{t}}}{\partial \mathbb{E}_{t-1}[g(Z_t)]} \tag{3}$$

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 prices 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*.

Using expectations, in contrast, 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. In this model, the marginal benefit of adaptation is the adaptation vector multiplied by the revenue value of those changes, denoted

$$B(\mathbf{A}_t) = \frac{\partial \mathbb{E}_{t-1}[y^*_t]}{\partial \mathbf{x}^*_t} \cdot \frac{\partial \mathbf{x}^*_t}{\partial \mathbb{E}_{t-1}[g(Z_t)]} \tag{4}$$

where arguments of the maximized output and choice variables have been suppressed for brevity. Estimating this value is one of the primary goals of the paper. 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 saves on such costs.

In the continuous model presented here, adaptation is welfare neutral on the margin.[21] For discrete adaptation like changes in land use or technology choice, Guo and Costello (2013) show that the benefits of small increases in adaptation can exceed the costs. For those cases, the derivatives in Equation (3) can be replaced by differences. Adaptation is then the change in inputs, broadly defined, in response

[21]In papers that use the envelope theorem to estimate climate effects, the equivalence between marginal benefits and marginal costs of adaptation at optimum is exploited to recover direct effect estimates (Deschênes and Greenstone, 2007, Hsiang, 2016).

to changes in the environment and the benefit of adaptation bounds the costs of adaptation from above.

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

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

Under the assumption that all adaptations are forward looking, the direct effect of weather on revenue is equal to the direct effect of weather on profit. This assumption rules out amelioration behavior which happens after the state realizes (Graff Zivin and Neidell, 2013). In a more general model, discussed in Section A, that incorporates choices made after the state realizes, it can be seen that both expectations and realizations of weather enter a more general adaptation term and that the method outlined here provides an upper bound on the direct effect.

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.[22] In the next section, I argue that weather forecasts can provide a good proxy for agent beliefs.

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

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

Thus, changes in realized weather identify the direct effect because

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

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

Showing that $\mathbb{E}_{t-1}[\partial y_t/\partial \mathbb{E}_{t-1}[g(Z_t)]] = \partial \mathbb{E}_{t-1}[y_t]/\partial \mathbb{E}_{t-1}[g(Z_t)]$ 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

[22] Parametric identification results with a known functional form for g are shown here. For the more general case with non-separable inputs and non-parametric identification, see Section A.1.

partial derivative of *ex ante* output with respect to expected weather.

2.2 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 would be

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

where $\mathbb{E}^p_{t-1}[g(Z_t)]$ is the private expectation that the agent holds about the weather next period and $g(z_t)$ is the realization of weather.

Observing these private expectations is usually not possible in practice, and finding good proxies for agent beliefs is challenging in general. Researchers studying weather effects, however, are well positioned to employ a method with many good theoretical properties: using professional forecasts of the relevant weather process as the measure of agent beliefs. 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 being observable.

Professional forecasts will provide a good measure of agent beliefs under the assumptions that: (1) the forecasts are public, (2) that 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 being proxied for (Wooldridge, 2010, ch.4). Redundant means that if the true expectations of the agent were observed, then the public forecast would not be helpful in explaining revenue. Formally, that $\mathbb{E}\left[y \middle| g(z), \mathbb{E}^p[g(Z)], \widehat{g(z)}\right] = \mathbb{E}\left[y \middle| g(z), \mathbb{E}^p[g(Z)]\right]$. 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

$$\mathbb{E}^p_{t-1}[g(Z_t)] = \theta_0 + \theta_1 \widehat{g(z_t)} + \xi_t \tag{7}$$

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

$$y_t = \alpha_0 + \alpha_2\theta_0 + \tilde{\alpha}_1 g(z_t) + \theta_1\alpha_2\widehat{g(z_t)} + \alpha_2\xi_t + \nu_t. \tag{8}$$

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 $\mathbb{E}[g(z_t)\xi_t] = 0$. 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, $\widehat{g(z_t)}$ will be an unbiased estimate of $\mathbb{E}^p_{t-1}[g(Z_t)]$ 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.

This discussion highlights two important benefits of including accurate measures of agent beliefs in an estimating equation for weather effects. First, excluding a measure of beliefs will lead to omitted variable bias. This bias is straightforward to sign because the agent's beliefs should be positively correlated with realizations of weather, so one only needs to know whether weather is positively or negatively correlated with revenue to determine whether the bias will be positive or negative. In the empirical setting, it will be shown that omitting forecasts leads to over-estimates of the direct effect (α_1 in the context of Equation (6)).

Second, an alternative approach to measuring agent expectations that is employed implicitly in much of the literature is to use average weather. 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.[23] In cases where the relevant stochastic

[23]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. In general, by measuring true beliefs with error, average weather will provide attenuated estimates of adaptation and exaggerated estimates of direct effects.

variable is stationary and agents have unchanging beliefs, then adaptation as defined by Equation (3) will be zero, and the appropriate way to study adaptation would be through changes in returns to, or prices for, adaptation mechanisms. On the other hand, using contemporary averages makes the assumption that agents have, and act on, perfect foresight 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 also assumes that the period over which weather is averaged is equal to the period over which beliefs about the weather are fixed.[24]

2.2.1 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 $\mathbb{E}[g(z_t)\xi_t] > 0$, so the usual omitted variable bias formula can be applied to find that $\operatorname{plim}|\tilde{\alpha}_1| = \left|\alpha_1 + \alpha_2 \frac{\mathbb{C}\mathrm{ov}(\xi, g(z))}{\mathbb{V}(g(z))}\right|$. 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 Stein (1956) and 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

[24] 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.

are partly disjoint, then one could see bias in α_2. For instance, if the firm consumes its 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 benefit of adaptation and direct effect of climate fluctuations on 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 price effects from climate variation.[25] The primary variable cost comes from diesel fuel, a globally traded and produced commodity. 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.

Albacore (*Thunnus alalunga*) typically stay in waters with sea surface temperature between 15 and 20°C (Childers et al., 2011). They also follow oceanic fronts with strong temperature gradients which limit the movement of their prey. The temperature preferences of albacore make them highly responsive to changes in climate.

[25] Table A7 shows that ENSO does not have a strong contemporaneous effect on wholesale albacore price. There is a small but significant increase in price from the one-year lag of ENSO.

The preferences of the albacore have led harvesters to develop rules of thumb based on sea surface temperature and ocean conditions including water color and clarity when determining where to try to catch fish (Clemens, 1961, Laurs et al., 1977, 1984, Childers et al., 2011).

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).[26] ENSO, therefore, generally entails more intensive and costly search for fish. In interviews, harvesters indicate that if uncertainty about optimal fishing location is too high or if expected fishing grounds are too distant from shore, 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 *detectable* prior to the deployment of the Tropical Atmosphere Ocean (TAO) array of weather buoys between 1984 and 1994 (Hayes et al., 1991).[27]

Skillful forecasts of ENSO were developed starting in the mid-1980s. Cane et al. (1986), a group of researchers at the Lamont-Doherty Earth Observatory (LDEO), published the first coupled ocean-atmosphere forecast of ENSO, termed 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 Di-

[26] 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 I find.

[27] 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).

agnostics 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.[28] By the end of 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 are a major improvement over simple persistence or autoregression-based predictions. The correlation between the three-month-ahead forecast and realized ENSO is 0.87 versus 0.74 for the correlation between ENSO and the three-month lag of ENSO. Regressing Niño 3.4 index time series 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 ENSO on the three-month lag, the coefficient is 0.74 and the R^2 is 0.55. Thus, the forecasts provide a 20 or 30% improvement, on average, over a simple lag-based forecast. Analyses of forecast accuracy and performance over time can be found in Barnston et al. (2010, 2012) as well as Figure A4. 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.[29]

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.[30]

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

[29]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).

[30]As discussed above, fishing trips generally last between 2 weeks and 1 month. As I discuss further below, the choice of horizon is also relatively unimportant if the researcher is only interested in separately identifying the benefit of adaptation and direct effect. Use of multiple forecast horizons

3.2 Dataset construction

For estimation, data on equatorial and North Pacific sea surface temperatures, ENSO forecasts, vessel-level fish catch, and relevant prices need to be combined. Here, I briefly describe each dataset used in the analysis. 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 temperature anomalies relative to a 30-year temperature average for a region of the equatorial Pacific Ocean. NOAA's Climate Prediction Center (CPC) publishes monthly average temperature anomalies in what is known as the Niño 3.4 region of the Pacific, a rectangular area ranging from 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 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 me with digital records for the period from

can be helpful for understanding the timing of adaptation.

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 but using data through May). 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.[31] The National Marine Fisheries Service (NMFS) requires the 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. Section B.5 provides details. 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, so albacore prices are only available for those landings (about 78% of the primary estimation sample). I perform my primary analysis on the sample where 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.[32]

The vessels in the sample use #2 marine diesel fuel. Where available, the price

[31] 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.

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

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, number 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 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.

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 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 will estimate linear specifications regressing revenue or output on the one-month lag of ENSO and its forecast. The lag between changes in ENSO in the equatorial Pacific and the effects being felt in the North Pacific suggests that the first lag of ENSO is what affects the fishery.

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, estimating the effect of ENSO on output in the period prior to the introduction of forecasts provides evidence on the correct functional form without needing to account for agent beliefs.[33] Figure 1 implements this test, showing

[33] As discussed in the background section, in the mid 1980s NOAA forecasters believed that skillful monthly horizon forecasts of monthly changes in ENSO were impossible. Therefore, the assumption that firms were unable to forecast deviations in ENSO relative to recent ENSO realizations might

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

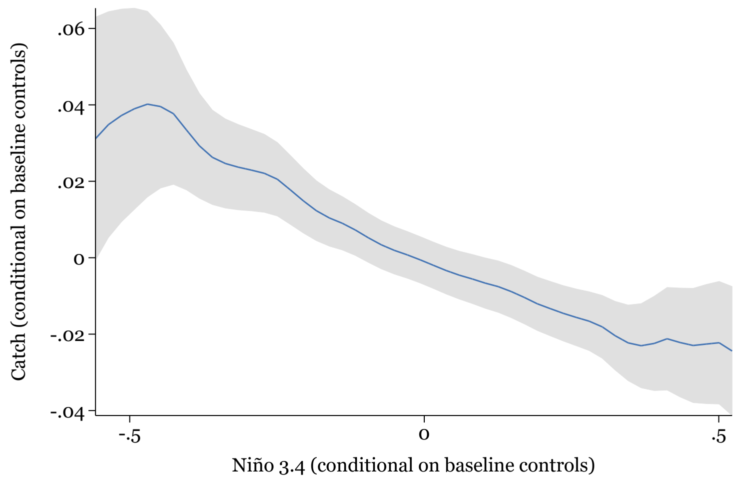

Notes: The figure shows a local linear regression (Epanechnikov kernel with bandwidth of 0.13) of monthly catch on the Niño 3.4 index the previous month. Both variables are residualized on month of year, year, and vessel fixed effects as well as two additional lags of the Niño 3.4 index. The Niño 3.4 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.

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. The figure shows that the relationship between ENSO and output in this period was linear.[34]

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

$$g(z_{t-1}) = \gamma_{\ell,0} + \gamma_{\ell,1} z_{t-1} \tag{9}$$

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 the fishery, then $\gamma_{\ell,1}$ will be negative. If cooler temperatures are harmful, then $\gamma_{\ell,1}$ will be positive.

---

be reasonable.

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

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

$$y_{it} = \beta_1 z_{t-1} + \beta_2 \hat{z}_{t-1} + \mathbf{z}'_{t-\ell}\alpha_z + \hat{\mathbf{z}}'_{t-\ell}\alpha_{\hat{z}} + \delta_{1,i} + \delta_{2,y(t)} + \delta_{3,m(t)} + \varepsilon_{it} \tag{10}$$

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 ENSO.[35] The variable ε is assumed to be a vessel and time-varying, stochastic error term.

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 $\hat{z}_{t-1}$, relative to that of the coefficient on z_{t-1}. The larger the magnitude of β_2 relative to β_1, 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 Results for ENSO effects and adaptation

5.1 Estimates of adaptation and direct effect

Table 2 shows results from implementing the primary identification strategy. Each column shows estimates of versions of Equation (10) 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

[35] 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). Given that the focus is on estimating the overall benefit of forward-looking adaptation, however, the exact horizon of the forecast is not crucial. Forecasts at different horizons are positively correlated, so they will recover related estimates. The crucial distinction is between forecasts and realizations.

each vessel. All dependent variables are standardized to have mean 0 and standard deviation of 1. 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 a radius of 30km. Autocorrelation in the errors is accounted for using 24 months of lags (Newey and West, 1987).[36]

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 (10) 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 effect of these variables is similar when using either output or revenue as the outcome. The theoretically

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

appropriate measure to use is revenue. The results indicate, however that ENSO does not have a strong effect on price.[37] Therefore, either measure can be interpreted following the conceptual model. Due to the presence of some interpolation in the revenue measure, I focus primarily on output throughout the rest of the paper.[38]

The first and third columns regress catch and revenue on measures of the realized strength of the one-month lag of the Niño 3.4 index but do not include forecasts. The results illustrate the omitted variable bias that occurs if forecasts are not included in the regression. The coefficients on realized Niño 3.4 indicate that ENSO has a moderate, 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 $\widehat{Niño\ 3.4}$).[39] The two coefficients in the table correspond to β_1 and β_2 from Equation (10). 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 in particular 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. Given the importance of accurate

[37] Table A7 shows that a 1 standard deviation increase in ENSO is associated with a 2% decrease in the wholesale price of albacore.

[38] However the effect of the imputation is assessed in robustness Table A1 and found to be small.

[39] The second and third lags of the forecast are also added as controls. All specifications include vessel, year, and month fixed effects as well as the second and third lag of the realized Niño 3.4 index.

estimates of climate damages for setting optimal policy, this omitted variable bias is substantial.

Continuing to interpret the estimates within the model, the effect of forecasts identifies the benefit of forward-looking adaptation. The large forecast coefficients indicate that forward-looking adaptation is a relatively more important driver of the effect of climate variation on the industry than are either *ex post* adaptation or the direct, realized effect. 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 are consistent with the concern many authors have raised about adaptation biasing the direct effect of weather.[40]

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.

The amount of forward-looking adaptation relative to the total effect gives a summary measure of the effectiveness of adaptation in this setting. In the context of the model, this value is $B(\mathbf{A})$ from Equation (4) (the marginal benefit of adaptation) divided by the total effect ($d\mathbb{E}[y^*]/d\mathbb{E}[g(Z)]$). I denote this value by $B_n(\mathbf{A})$ because it is the normalized benefit of adaptation.

With a linear specification, $B_n(\mathbf{A})$ can be calculated using estimates of $\frac{\beta_2}{\beta_1+\beta_2}$. For output, $B_n(\mathbf{A})$ is 0.75 (95% confidence interval of 0.57 to 0.93), indicating that 75% of the total effect of ENSO on output is due to adaptation. For revenue, $B_n(\mathbf{A})$ is

[40] 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.

0.60 (95% confidence interval of 0.44 to 0.76).[41] In both cases, adaptation is well above zero.

These estimates show that climate variation has a strong effect on the firms in this fishery. But, the effect is substantially overstated if forecasts are not included in the regression. Interpreting these forecasts as measures of the expectations held by the firms, the results imply that forward-looking adaptation makes up the bulk of the effect of climate variation on the industry.

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 fall 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 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 autocor-

[41] Standard errors calculated using the delta method.

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)
$\widehat{\text{Niño 3.4}}$	-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 (10) 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.

relation 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 is collapsed to the monthly level. The spatial standard errors are preferred for the baseline specification, however, because ENSO does have local effects on fishing conditions that vary smoothly over space (see Appendix Figure A5), so year-month clustering is likely to be too conservative.

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

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)
$\widehat{\text{Niño 3.4}}$	-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 (10) 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.

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.

Two additional results, not reported in the tables, check multiplicative separability and sensitivity to outliers. Multiplicative separability is tested by including the interaction between Niño 3.4 and the forecast in the primary specification. The coefficient on the interaction is insignificant but practically close to zero. Second, the primary specification is run using output that has been Winsorized at the 1% level. The logbook records are of high quality, but Winsorization should reduce the effect of any potential outliers. Inference is unchanged in this case.

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.

6 Adaptation mechanisms

This section explores how the firms in this fishery 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. One of the benefits of the methodology in this paper is that the researcher need not know about or have data on all potential adaptation mechanisms to still gain an understanding of the benefit of adaptation. 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 is filled quickly, then fishing hours would go down. If fishing is poor, the crew may continue to fish longer to make up for the shortfall or may stop fishing earlier to change fishing locations. In response to anticipated ENSO, harvesters decrease their hours fished per day by just over 5%. Whether this is due to poor fishing conditions or not, this

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 (10) 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.

reduction in hours worked per day entails a reduction in intensive-margin effort.

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. 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.[42] 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 if ENSO is anticipated than if it is not. If harvesters are using fewer lines to save on labor costs, this could represent an important overall reduction in variable costs.

Another major source of variable cost for the vessels is the burning of fuel during transit and fishing. Table 5 Column 3 shows the effect of ENSO on vessel movement. Section B.4 provides details on this measure, but the basic method is to use the latitude and longitude records each day to calculate day-to-day movement. Such a calculation will miss intra-day movement. The results indicate that harvesters move less if they anticipate worse ENSO conditions. In fact, expecting a moderate ENSO

[42] This is in contrast to methods like long-lining where a single, long fishing line might have hundreds or thousands of baited hooks.

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 (10) 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

The main results from Table 2 incorporate adaptation that is occurring both between fishing trips and once a vessel is out fishing. 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 fishing* 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.

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.

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)
$\widehat{\text{Niño 3.4}}$	-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 (10) 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.

6.4 Net Revenue

Regressing forecasts and realizations of ENSO on output and revenue, as is done in the previous section, 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, a method employed recently by Carleton et al. (2018).[43] 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 or the wages paid

[43] 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, Lemoine and Traeger, 2014).

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 (10) 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.

to crew. The one input that can be consistently calculated is movement during fishing trips. To arrive at movement costs, I multiply movement by the average real price of fuel, based on port-level surveys. 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 10 to 20% of the variable cost of running an albacore vessel, so the resulting costs are scaled to constitute 20% 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 primary estimation sample (where net revenue and fuel costs are observed). As expected from the movement results in Section 6.2, fuel costs decline substantially when ENSO is anticipated. 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.[44] 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 and Udry (2014a) 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

$$y_{it} = \beta_0 + \beta_1 z_{t-1} + \beta_2 \hat{z}_{t-1} + \beta_3 \hat{z}_{t-1}\text{qual}_{t-1} + \mathbf{x}'_{it}\alpha + \varepsilon_{RA,it} \tag{11}$$

where the new variable *qual* is a measure of the *ex ante* quality of the forecast. 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 modified in two ways (Hamill and Juras, 2006). 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.

[44] These changes in net revenue are due to changes in firm behavior rather than through changes in albacore or 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.

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 (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.

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 relative to the direct effect 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 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.[45]

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 normalized benefit of adaptation for novice, experienced, and highly experienced vessels.[46] For a novice vessel (25$^{\text{th}}$ percentile experience), adaptation is about 20% lower than for a very experienced vessel (75$^{\text{th}}$ percentile experience).

These results raise the question of whether the high adaptability of more experienced firms can be replicated by less experienced firms, potentially through training, mentorship, or some other mechanism. Together with the previous results on risk aversion, they suggest that forecasts will be more effective at promoting adaptation if they are high quality and distributed to knowledgeable firms.

8 Conclusion

Environmental impacts from a variety of source are currently large and, for many important cases, are not being address by collective action at a scale appropriate to the potential damages. If public policy is not appropriately aggressive, then individual and firm adaptation will need to play an outsize role in damage reduction. Adaptation does not occur in a vacuum, however. Individuals need to know about their own risks

[45] See Kala (2015) for further evidence on learning about forecastable weather in an agricultural context.

[46] The experience trends are defined at the vessel level because captain identities are not consistently reported in the logbook records.

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)
$B_n(A)$ low experience	0.61***	0.70***	0.55***
	(0.076)	(0.083)	(0.074)
$B_n(A)$ medium experience	0.76***	0.79***	0.73***
	(0.080)	(0.087)	(0.078)
$B_n(A)$ 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 (10) 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.

to make informed choices over potential adaptive responses. The importance of this issue makes it crucial to assess the role of information in facilitating adaptation. The effect of information on adaptation also allows one to use informational changes to estimate the effect of this adaptation.

In the setting of one large driver of global climate—ENSO—and firms with flexible production, this paper assesses the degree of forward-looking adaptation using an estimating equation informed by a simple structural model of adaptation to a stochastic weather process. 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 allows harvesters to take action that substantially reduces the direct effects of this climate variable.

From a methodological standpoint, the empirical strategy presented here has the potential to be applied to many settings. The novel collection of ENSO forecasts assembled for the project and the estimation strategy 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. Applying the methodology—either to understand ENSO or for weather effects more broadly—can help clarify the extent of adaptation and can improve estimates of the effect of weather conditional on that adaptation.

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 (IPCC, 2014). 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. Moreover, as Hornbeck and Keskin (2014) shows empirically, long-run adaptation can be perverse in the sense that a relaxation of one constraint can allow individuals or firms to place themselves in an even more climate-exposed long-run position. 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 here point to the direct value of policy-driven information provision. Information externalities imply that public provision of forecasts of weather and climate changes can have a positive welfare impact even if adaptation mechanisms themselves are private.

Outside the context of environmental adaptation, the results demonstrate that analysis of forecasts of environmental processes can help improve understanding of questions in firm and consumer responses to information. For instance, the theory of adaptation shares a formal similarity with theories of firm flexibility introduced by Stigler (1939). Such theories are generally difficult to test due to a lack of data on expectations. Using environmental forecasts can allow for investigation of trade-offs in stochastic settings, an area ripe for further development.[47]

[47] For example Downey et al. (2021) investigates *ex ante* responses to ENSO-driven rain forecasts to study costly labor adjustment.

References

Albacore Working Group (2014). Stock assessment of albacore tuna in the north pacific ocean in 2014. Technical Report August.

Allen, R., J. Graff Zivin, J. Shrader, J. Zivin, and J. Shrader (2016). Forecasting in the presence of expectations. *The European Physical Journal Special Topics 225*(3), 539–550.

Auffhammer, M. (2018a). Climate adaptive response estimation: Short and long run impacts of climate change on residential electricity and natural gas consumption using big data. *NBER Working Paper No. 24397*, 1–45.

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

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.

Bento, A., N. S. Miller, M. Mookerjee, and E. R. Severnini (2020). A unifying approach to measuring climate change impacts and adaptation. *NBER Working Paper No. 27247*, 49.

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., M. Craxton, C. D. Kolstad, C. Onda, H. Allcott, E. Baker, L. Barrage, R. Carson, K. Gillingham, J. Graff-Zivin, M. Greenstone, S. Hallegatte, W. M. Hanemann, G. Heal, S. Hsiang, B. Jones, D. L. Kelly, R. Kopp, M. Kotchen,

R. Mendelsohn, K. Meng, G. Metcalf, J. Moreno-Cruz, R. Pindyck, S. Rose, I. Rudik, J. Stock, and R. S. J. Tol (2016). Opportunities for advances in climate change economics. *Science 352*(6283).

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

Butler, E. E. and P. Huybers (2013, November). Adaptation of us maize to temperature variations. *Nature Climate Change 3*(1), 68–72.

Cane, M., S. Zebiak, and S. Dolan (1986). Experimental forecasts of el niño. *Nature 321*, 827–832.

Carleton, T., M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, A. Jina, R. E. Kopp, K. McCusker, I. Nath, et al. (2018). Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits.

Carleton, T., A. Jina, M. T. Delgado, M. Greenstone, T. Houser, S. M. Hsiang, A. Hultgren, R. E. Kopp, K. E. McCusker, I. B. Nath, J. Rising, A. Rode, H. K. Seo, A. Viaene, J. Yuan, and A. T. Zhang (2020). Valuing the global mortality consequences of climate change accounting for adaptation costs and benefits. *NBER Working Paper No. 27599*, 56.

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

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

Childers, J., S. Snyder, and S. Kohin (2011). Migration and behavior of juvenile north pacific albacore (thunnus alalunga). *Fisheries Oceanography 20*(3), 157–173.

Clemens, H. B. (1961). The migration, age, and growth of pacific albacore (thunnus germo), 1951-1958. *Fish Bulletin 115*, 128.

Colmer, J. (2021). Temperature, labor reallocation and industrial production: evidence from india. pp. 1–28.

Conley, T. (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. Jones, and B. Olken (2009). Temperature and income: Reconciling new cross-sectional and panel estimates. *American Economic Review: Papers & Proceedings 99*(2), 198–204.

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

Dell, M., B. Jones, and B. 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.

Deschênes, O., M. Greenstone, and J. S. Shapiro (2017). Defensive investments and the demand for air quality: Evidence from the nox budget program. *American Economic Review 107*(10), 2958–2989.

Downey, M., N. Lind, and J. G. Shrader (2021). Adjusting to rain before it falls. *CEEP Working Paper 17*, 1–38.

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. *El Niño and the Southern Oscillation: Multiscale Variability and Global and Regional Impacts.*, 89–112.

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

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

Fisher, A., W. Hanemann, M. Roberts, and W. Schlenker (2012). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: comment. *American Economic Review*.

Fried, S. (2021). Seawalls and stilts: A quantitative macro study of climate adaptation. *Federal Reserve Bank of San Francisco, Working Paper Series 2021-07*, 1–74.

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

Graff Zivin, J. and M. Neidell (2012). The impact of pollution on worker productivity. *American Economic Review 102*(7), 3652–3673.

Graff Zivin, J. and M. Neidell (2013). Environment, health, and human capital. *Journal of Economic Literature 51*(3), 689–730.

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

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.

Haaland, I., C. Roth, and J. Wohlfart (2021). Designing information provision experiments. *Journal of Economic Literature Forthcomin*, 1–77.

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.

Hayes, S. P., L. J. Mangum, J. Picaut, A. Sumi, and K. Takeuchi (1991). Toga-tao: A moored array for real-time measurements in the tropical pacific ocean. *Bulletin of the American Meteorological Society 72*(3), 339–347.

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. (2016). Climate econometrics. *Annual Review of Resource Economics 8*(July), 43–75.

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.

Imbens, G. W. and W. K. Newey (2009). Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica 77*(5), 1481–1512.

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.

Kala, N. (2015). Ambiguity aversion and learning in a changing world: The potential effects of climate change from indian agriculture. *Working Paper*, 1–54.

Kamenica, E. (2017). Information economics. *Journal of Political Economy 125*(6), 1885–1890.

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

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

Laurs, R., P. C. Fiedler, and D. R. Montgomery (1984). Albacore tuna catch distributions relative to environmental features observed from satellites. *Deep Sea Research Part A. Oceanographic Research Papers 31*(9), 1085–1099.

Laurs, R. M., H. S. Yuen, and J. H. Johnson (1977). Small-scale movements of albacore, thunnus-alalunga, in relation to ocean features as indicated by ultrasonic tracking and oceanographic sampling. *Fishery Bulletin 75*(2), 347–355.

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*, 483–494.

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

Lemoine, D. and C. Traeger (2014). Watch your step: Optimal policy in a tipping climate. *American Economic Journal: Economic Policy 6*(1), 137–166.

Massetti, E. and R. Mendelsohn (2018). Measuring climate adaptation: Methods and evidence. *Review of Environmental Economics and Policy 12*(2), 324–341.

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. (2000). Efficient adaptation to climate change. *Climate Change 45*, 583–600.

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

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.

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.

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

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

Repetto, R. (2008). The climate crisis and the adaptation myth. *Forestry and Environmental Studies Publications Series 9*(13), 1–24.

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

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

Rosenzweig, M. R. and C. Udry (2014b). Rainfall forecasts, weather, and wages over the agricultural production cycle. *American Economic Review Papers & Proceedings 104*(5), 278–283.

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

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.

Stein, C. (1956). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. *Proceedings of the Third Berkeley symposium on mathematical statistics and probability 1*(399), 197–206.

Stigler, G. (1939). Production and distribution in the short run. *Journal of Political Economy 47*(3), 305–327.

Taraz, V. (2017). Adaptation to climate change: Historical evidence from the indian monsoon. *Environment and Development Economics 22*(5), 517–545.

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

Toth, Z. and R. Buizza (2019). Weather forecasting: What sets the forecast skill horizon? In *Sub-Seasonal to Seasonal Prediction*, pp. 17–45. Elsevier.

Trenberth, K. (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. Technical report.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data* (2 ed.), Volume 58.

Appendix for online publication

A Model extensions

A.1 Non-separable weather

The model in Section 2 assumed that weather and inputs were multiplicatively separable. Without assuming this separability, the definition of adaptation and estimation strategy still hold, but the relatively simple dependence of adaptation on a single function of weather will no longer hold.

Consider a single input model but without the separability assumption. Formally, let the firm solve

$$\max_x \mathbb{E}_{t-1}[\pi_{it}] = p_{1t}\mathbb{E}_{t-1}[f(x_{it}, Z_{it})] - p_{2t}x_{it}. \tag{12}$$

Suppressing entity and time subscripts and using subscripts on equations to denote partial derivatives, the first order condition will be

$$p_1\mathbb{E}[f_1(x, Z)] - p_2 = 0 \tag{13}$$

By the implicit function theorem, one can find

$$\frac{\partial x}{\partial \mathbb{E}[Z]} = -\frac{\partial \mathbb{E}[f_1(x, Z)]}{\partial \mathbb{E}[Z]}\left(\frac{\partial \mathbb{E}[f_1(x, Z)]}{\partial x}\right)^{-1} = -\frac{\frac{\partial \mathbb{E}[f_1(x,Z)]}{\partial \mathbb{E}[Z]}}{\mathbb{E}[f_{11}(x, Z)]}. \tag{14}$$

Similar expressions can be derived for other moments of the weather distribution, suggesting that a semiparametric procedure for estimating this more general model would be to include progressively higher moments of the weather forecast distribution in the estimating equation. Such a procedure would require a rich forecast (of the probability density, for instance) or a simple weather process. Formal identification of this model comes from application of recent results in identification of nonparametric instrumental variables models with non-separable error.

Let the optimal input choice be

$$x_t^* = \underset{x}{\text{argmax}}\{p_{1t}\mathbb{E}[f(x_t, Z_t)|\hat{\mathbf{z}}_{t|t-1}] - p_{2t}x_t\}, \tag{15}$$

where $\hat{\mathbf{z}}_{t|t-1}$ is the vector of forecasts of moments of the distribution of Z_t that the agent forms based on the information set $\mathscr{G}_{t-1}$.[48] This problem yields an optimal choice for x denoted $x_t^* = h(\hat{\mathbf{z}}_{t|t-1}, \eta_t)$ where η contains everything that shifts factor

[48] Under loss functions discussed in Section A.4, this vector is simply the conditional expectation of Z_t.

demand other than expectations about the weather. Finally, denote deviations from expected weather by $\varepsilon_{n,t} = \mathbb{E}[Z_t^n] - \hat{z}_{n,t|t-1}$, where n indexes the moments of the weather distribution, and collect these deviations in the vector ε_t.

Assuming that x_t is strictly monotonic in η_t and that $\hat{\mathbf{z}}_{t|t-1}$ is independent of η_t and ε_t, the results from Imbens and Newey (2009) can be applied to identify f. Two of these assumptions are natural in this setting. In the model, η contains prices, so the law of demand gives monotonicity. A sophisticated forecaster will ensure that $\hat{\mathbf{z}}$ is exogenous with respect to ε_t. Finally, a maintained assumption is that prices are independent of expected weather, leading to independence of η and $\hat{\mathbf{z}}$.

This more general identification reinforces the intuition from the separable case presented in the body of the paper. Forecasts errors are useful for identifying direct effects of weather, and under the assumption that forecasts only affect inputs, the factor demand can be fully recovered even if prices are not observed.

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

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

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,

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

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

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

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 $\mathbb{E}[g(Z)]$ 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

$$\max_{x_{2t}} \pi_t = f(x_{1t}^*, x_{2t})g(z_t) - p_1 x_{1t}^* - p_2 x_{2t} \tag{16}$$

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

$$E[L^p(Z_t, \hat{Z}_t, \mathbf{X}(\hat{\mathbf{Z}})_t, \mathbf{p}_t)|\mathscr{G}_{t-h}] \tag{17}$$

where we now allow inputs to be a vector and expectations about the future weather are denoted by $\hat{Z}$. $\mathscr{G}_t \in \mathbb{F}$ 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 (17) 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, $L_1^p(Z_t, \hat{Z}_t, \mathbf{X}_t, \mathbf{p}_t)$, 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 $Z_t - \mathbb{E}[Z_t|\mathscr{G}_{t-h}]$.

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 $s^p_{t|t-h} = \mathbb{E}[z_t|\mathscr{G}_{t-h}]$ (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 $\mathscr{F}_{t-h}$

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

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

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

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: $\mathbb{C}\text{ov}(s_{t+h|t}, x_t) = 0$ for all $x_t \in \mathscr{F}_t$

3. Longer lead forecasts are less precise:

    - $\mathbb{V}(s_{t+h|t}) \leq \mathbb{V}(s_{t+H|t})$ for all $h \leq H$
    - $\mathbb{V}(\varepsilon_{t+h|t}) \leq \mathbb{V}(\varepsilon_{t+H|t})$ for all $h \leq H$ where $\varepsilon_{t+h|t} = z_{t+h} - s_{t+h|t}$ 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 $\mathscr{G}_t \supseteq \mathscr{F}_t$ and $(\mathscr{F}_t)_{t\geq 0}$ is strictly monotonic, then there exists a forecast $s_{\tau|t+k}$ such that $\mathbb{V}(\varepsilon_{\tau|t+k}) = \mathbb{V}(\varepsilon^p_{\tau|t})$ for $k \geq 0$.*

Proof. Forecast properties gives us that $\mathbb{V}(\varepsilon_{\tau|t}) \geq \mathbb{V}(\varepsilon^p_{\tau|t}) \geq \mathbb{V}(\varepsilon_{\tau|\tau})$.

Therefore, by continuity there must exist a $k \geq 0$ 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 $\mathscr{G}_t \subseteq \mathscr{F}_t$. 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 $\mathscr{G}_t \subseteq \mathscr{F}_t$, then $s^p_{t+h|t} = s_{t+h|t}$.*

Proof. The Granger conditions imply that $s^p_{t+h|t} = \mathbb{E}[z_{t+h}|\mathscr{G}_t]$, so by Lemma A.1 and MSE-forecast property 3, $\mathscr{G}_t \subseteq \mathscr{F}_t$ implies

$$\mathbb{V}(\varepsilon^p_{t+h|t}) \geq \mathbb{V}(\varepsilon_{t+h|t})$$

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: $\mathscr{G}_t \nsubseteq \mathscr{F}_t$

To estimate adaptation, we are interested in $\frac{dy}{ds^p}$. If we observed s^p and $\mathscr{G}_t \supseteq \mathscr{F}_t$, 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 $\mathscr{G}$ and $\mathscr{F}$, 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, $\mathscr{G}$ 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 $\mathscr{G}_t \not\supseteq \mathscr{F}_t$ and $\mathscr{G}_t \not\subseteq \mathscr{F}_t$. Here, because forecasts based on $\mathscr{F}_t$ are public, the firm will incorporate the public forecast into their private information, leading to $\tilde{s}^p_{t|\tau} = g(s^p_{t|\tau}, s_{t|\tau})$. 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

$$\begin{aligned}\tilde{s}^p_{t|\tau} &= \frac{(w^p s^p_{t|\tau} + w s_{t|\tau})}{w^p + w} \\ \Rightarrow \frac{\partial \tilde{s}^p}{\partial s} &= \frac{w}{w^p + w}\end{aligned}$$

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.[49] 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

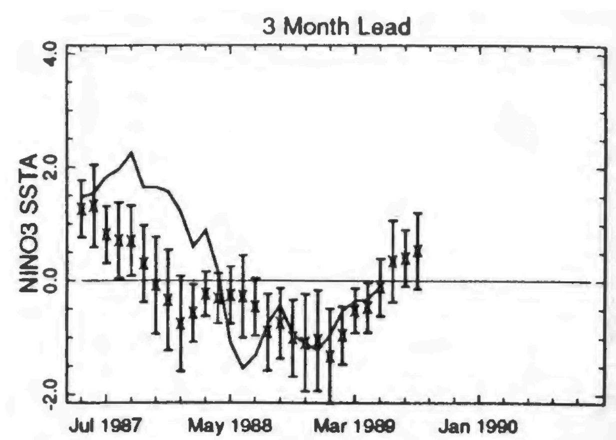

Notes: The figure shows an ENSO forecast issued in the Climate Diagnostics Bulletin in June of 1989. This figure is typical of the forecasts published between 1989 and 2002. The solid line shows the Niño 3 sea surface temperature anomalies and the X are forecasts (and back-casts). Whiskers are the historical 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

[49]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

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

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 $\gamma_{\text{month}}\text{gas}_t$ 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 (10) 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 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.

$$g(z_{t-1}) = \gamma_{q,0} + \gamma_{q,1} z_{t-1} - \gamma_{q,2} z_{t-1}^2 \tag{18}$$

With this function of weather, if agents are forming distributional beliefs about ENSO, then the correct forecast term to include would be $\widehat{g(z_{t-1})} = \gamma_{q,0} + \gamma_{q,1}\mathbb{E}_{t-h}[Z_{t-1}] - \gamma_{q,2}\mathbb{E}_{t-h}[Z_{t-1}^2]$, 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

$$\widehat{g(z_{t-1})} = \gamma_{q,0} + \gamma_{q,1}\mathbb{E}_{t-h}[Z_{t-1}] - \gamma_{q,2}\mathbb{E}_{t-h}[Z_{t-1}]^2 \tag{19}$$

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

$$\mathbb{E}[g(Z)] = \gamma_{q,0} + \gamma_{q,1}\mathbb{E}_{t-h}[Z_{t-1}] - \gamma_{q,2}\mathbb{E}_{t-h}[Z_{t-1}^2] \tag{20}$$

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

$$\mathbb{E}[g(Z)] - g(\mathbb{E}[Z]) = \gamma_{q,2}(\mathbb{E}_{t-h}[Z_{t-1}]^2 - \mathbb{E}_{t-h}[Z_{t-1}^2]) = \gamma_{q,2}\mathbb{V}_{t-h}(Z_t) \tag{21}$$

If one assumes that Z_t has constant variance over time, then (21) 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

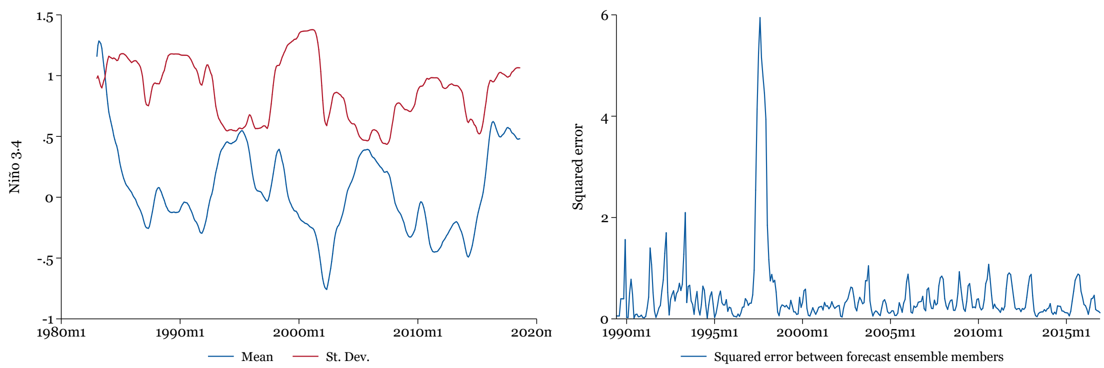

(a) ENSO Rolling St. Dev. and Average (b) Disagreement in Ensemble Members

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

$$y_{it} = \beta_{q,0} + \beta_{q,1} z_{t-1} + \beta_{q,2} z_{t-1}^2 + \beta_{q,3} \hat{z}_{t-1} + \beta_{q,4} \hat{z}_{t-1}^2 + \mathbf{x}_{it}' \alpha_q + \varepsilon_{q,it} \tag{22}$$

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 Sections B.8 and B.9.

B.8 Nonlinear effects of ENSO

Table A2 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 (22). Columns 2 and 4 add interactions between the forecast and realization of ENSO. For ease of interpretation, Table A3 shows the marginal effects for each model when both the forecast and realization of ENSO are equal to 1 (moderate El Niño).

Table A2: 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)
$\widehat{\text{Niño 3.4}}$	-0.19***	-0.19***	-0.17***	-0.18***
	(0.036)	(0.037)	(0.034)	(0.036)
$\widehat{\text{Niño 3.4}} \times \widehat{\text{Niño 3.4}}$	-0.088***	-0.18***	-0.076***	-0.22***
	(0.022)	(0.027)	(0.020)	(0.026)
Niño 3.4 × $\widehat{\text{Niño 3.4}}$		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 (22) 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.

The quadratic estimates reinforce the primary results from Table 2. First, in 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.[50] 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 A3: 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 A2. Standard errors calculated using the delta method. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

The normalized benefit of adaptation in these models depends on realizations of ENSO and the relevant forecast. Section B.9 discusses this measure in detail, but the basic intuition can be captured by comparing the marginal effect of forecasts to the sum of the marginal effect of forecasts and realizations. This value is the normalized marginal benefit of adaptation when a moderate El Niño hits. For output, the ratio is 0.86 for the quadratic specification in column 1 and 0.92 for the quadratic specification with interactions in column 2. The comparable measures for revenue are 0.71 and 0.83. In all cases, adaptation is a large fraction of the total effect of ENSO. The more complete analysis in Section B.9 reinforces this point, showing that normalized adaptation is high—greater than 50% for all methods of calculating the measure and near 100% in all cases for output.

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

[50]These results are reported in Section B.8.

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.9 Univariate measure of adaptation value for quadratic specifications

Comparing the value of adaptation with the residual, direct effect helps to determine whether the magnitude of total adaptation is large and aids in comparisons with other studies. In particular, the value of adaptation can be normalized by dividing by the total derivative of output with respect to a change in climate,

$$B_n(\mathbf{A}) = \frac{B(\mathbf{A})}{d\mathbb{E}_{t-1}[y_t^*]/d\mathbb{E}_{t-1}[g(Z_t)]}. \tag{23}$$

The normalization creates an intuitive adaptation index because the total change in output with respect to a change in climate can be decomposed into the change due to adaptation and the change due to direct effects.

$$\frac{d\mathbb{E}_{t-1}[y_t^*]}{d\mathbb{E}_{t-1}[g(Z_t)]} = \frac{\partial\mathbb{E}_{t-1}[y_t^*]}{\partial\mathbf{x}_t^*} \cdot \frac{\partial\mathbf{x}_t^*}{\partial\mathbb{E}_{t-1}[g(Z_t)]} + \frac{\partial\mathbb{E}_{t-1}[y_t^*]}{\partial\mathbb{E}_{t-1}[g(Z_t)]} \tag{24}$$

If the value of adaptation is high relative to the direct effect, then this value will be close to one. If adaptation is zero, this term will be equal to zero. The normalized benefit of adaptation also has a welfare interpretation under the assumption of continuous inputs. Given a choice over two continuous production technologies with the same costs, a firm would rather choose the technology with lower $\frac{\partial\mathbb{E}_{t-1}[y_t^*]}{\partial\mathbb{E}_{t-1}[g(Z_t)]}$ relative to $\frac{\partial\mathbb{E}_{t-1}[y_t^*]}{\partial\mathbf{x}_t^*} \cdot \frac{\partial\mathbf{x}_t^*}{\partial\mathbb{E}_{t-1}[g(Z_t)]}$, because the second term will be zero according to the first order condition and is therefore profit neutral, while the direct effect influences profit. Therefore, the firm would prefer the technology with a higher normalized benefit of adaptation, all else equal.

Estimating the normalized benefit of adaptation is straightforward if the effect of climate on revenue (and profit) is linear. In nonlinear specifications discussed in Section B.7, the calculation poses a problem, however, because the derivative of g will be zero at the peak of the quadratic curve. This will cause the mean of the total effect to be zero at this point, leading to division by zero. Figure 1 and the estimates from Table 2 show that the peak of the quadratic occurs near the center of the ENSO distribution, so this issue is a problem in practice.

There are a number of possible solutions to the division-by-zero problem, and in this section, I pursue three of them to compare their effect on the estimated, normalized benefit of adaptation. First, for the parametric specification used in the

Table A4: Normalized benefit of adaptation, quadratic models

Estimator of B_n(A)	(1) Catch Quadratic	(2) Catch Interaction	(3) Revenue Quadratic	(4) Revenue Interaction
Median	0.74	0.75	0.64	0.67
	[0.68,0.80]	[0.70,0.80]	[0.59,0.69]	[0.63,0.71]
Conditional average	0.71	0.77	0.61	0.51
	[0.38,1.14]	[-0.06,1.60]	[0.44,0.81]	[0.05,1.15]
Limit as Niño 3.4 $\rightarrow \infty$	0.71	0.87	0.72	1.04
	[0.53,0.88]	[0.71,1.02]	[0.53,0.90]	[0.82,1.26]

Notes: The table shows results from three estimators of Equation (23) using monthly data and quadratic specifications. The dependent variable in each column corresponds to a model from Table A2. 95% confidence intervals are shown in parentheses and are calculated by the delta method for the limit and by bootstrap in the case of the conditional mean and the median.

baseline results, one can take a limit of the numerator and denominator of Equation (23) as Niño 3.4 goes to infinity. In the quadratic specification without interactions, this limit is not a function of ENSO, and $B_n(A)$ simplifies to be $\beta_4/(\beta_2+\beta_4)$, where the coefficients are those from Equation (22). This method has the advantage that standard errors can be easily calculated using the delta method under the assumption that $\beta_4 \neq 0$. When including interactions between Niño 3.4 and the forecast, the limit becomes

$$B_n(A) = \frac{2\beta_4 + \beta_5}{2\beta_2 + 2\beta_4 + 2\beta_5}$$

where β_5 is the coefficient on the interaction between Niño 3.4 and the forecast.

Second, one can calculate the median of $B_n(A)$ using the empirical distribution of ENSO. The median is less subject to outliers caused by division by zero. For both the conditional mean and the median, standard errors are calculated by bootstrap over the parameter estimates from Table 2 and the empirical distribution of ENSO given by Niño 3.4 values from 1989 to 2010. Results using 3,000 bootstrap replications are shown.

Third, one can condition on being away from the point of zero slope when estimating the expectations in Equation (23). This method is convenient, but it also has a nice interpretation as capturing adaptation to deviations from conditions to which the firm is well adapted. In the table, I condition on the Niño 3.4 index being greater than 0.5 of less than -0.5 away from the singularity. The results end up being almost identical to the median approach. The 95% confidence interval is again based on 1,000 bootstrap replications.

In all cases, total adaptation is clearly statistically different from 0, in contrast to recent studies of adaptation in other settings (Burke and Emerick, 2016, Dell et al., 2012, Schlenker et al., 2013). In fact, the point estimates are consistently greater than one half. The results show that forward-looking adaptation is substantial in this setting.

Three potential sources of bias also suggest that, if anything, these estimates understate total adaptation. First, if harvesters have private information about ENSO that is not captured by the public forecasts, then the model in Section 2 shows that estimated, forward-looking adaptation will be attenuated. Second, if some adaptation mechanisms can occur after the effects of ENSO events are known, then forward-looking adaptation is only part of the total adaptation response, and part of the direct effect would actually be an *ex post* adaptive response. I find some evidence for *ex post* adaptation in Section 6, but the small magnitude of the realized ENSO coefficients in Table 2 allows one to infer that there is, at most, only limited adaptation of this type. Third, because the pre-2002 forecasts had to be digitized from printed records, some (likely classical) measurement error probably exists. The ENSO index is consistently well measured over the estimation sample period since it occurs after the advent of satellite and buoy measurement, so the measurement error in the forecasts should lead to attenuation of the forecast coefficient.

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 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 in this setting. 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)
$\widehat{\text{Niño 3}}.4_{t-1}$		-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 (11) 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-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 (11) 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

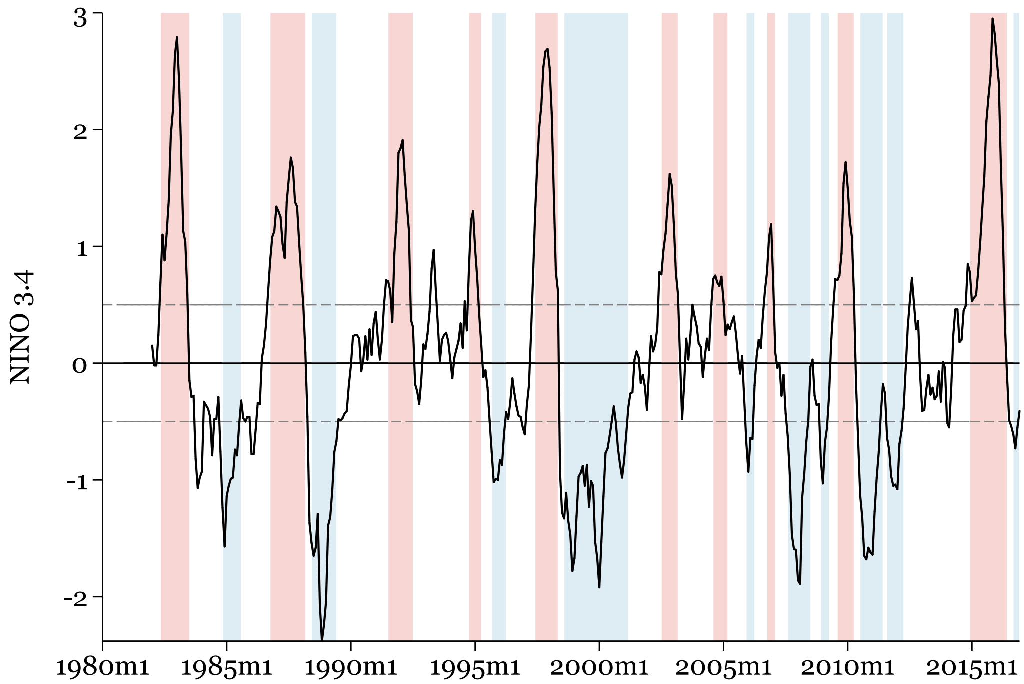

Notes: The ENSO cycle is measured here by the Niño 3.4 index, which is the three month moving average of sea surface temperature anomalies from the Niño 3.4 region of the equatorial Pacific Ocean. Values above 0.5 indicate an El Niño and values below -0.5 indicate La Niña, as denoted by the red and blue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

Figure A4: ENSO Forecast Skill

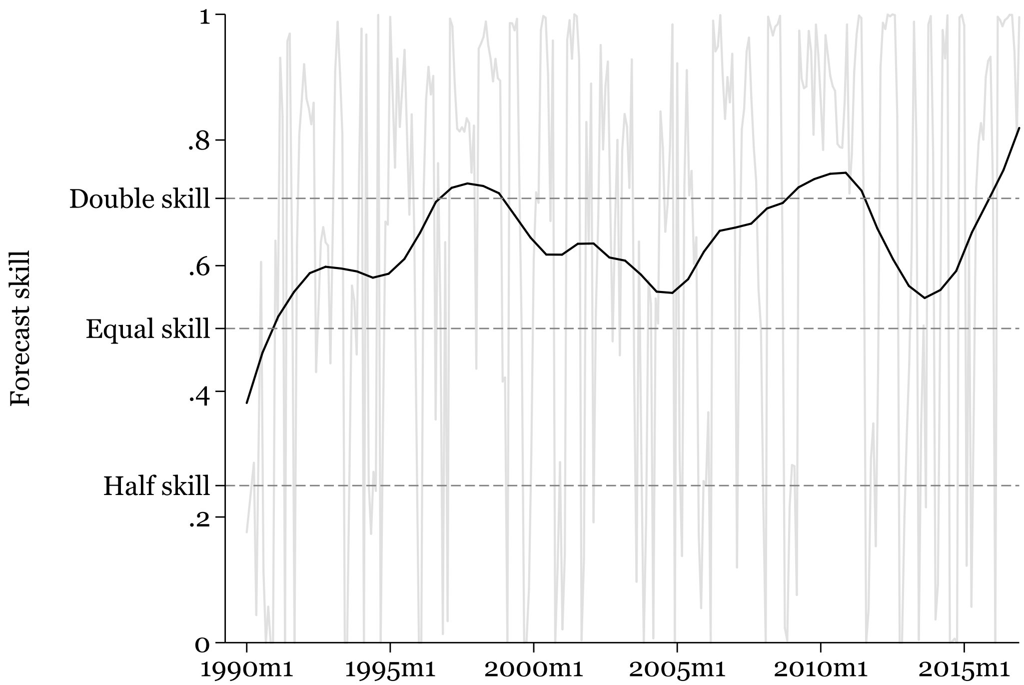

Notes: Forecast skill measured by a normalized version of the Brier skill score is indicated by the light gray line. Skill is the exponential of log of 0.5 times squared error of the forecast divided by the squared error of a naïve persistence forecast. The moving average of monthly skill is given by 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

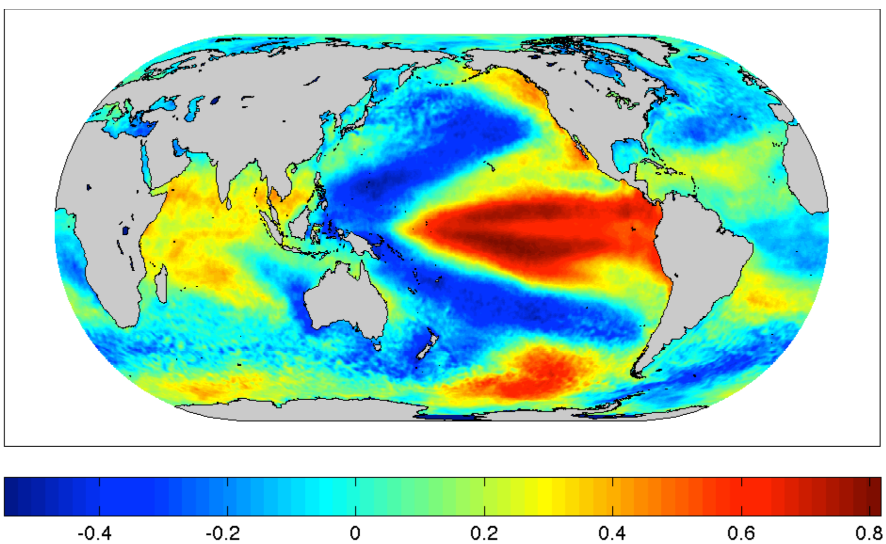

Notes: The heat map shows correlation between the one month lag of the Niño 3.4 index and sea surface temperature for each quarter degree latitude-longitude grid cell.

Figure A6: Fishing and Transiting Locations for Daily Observations

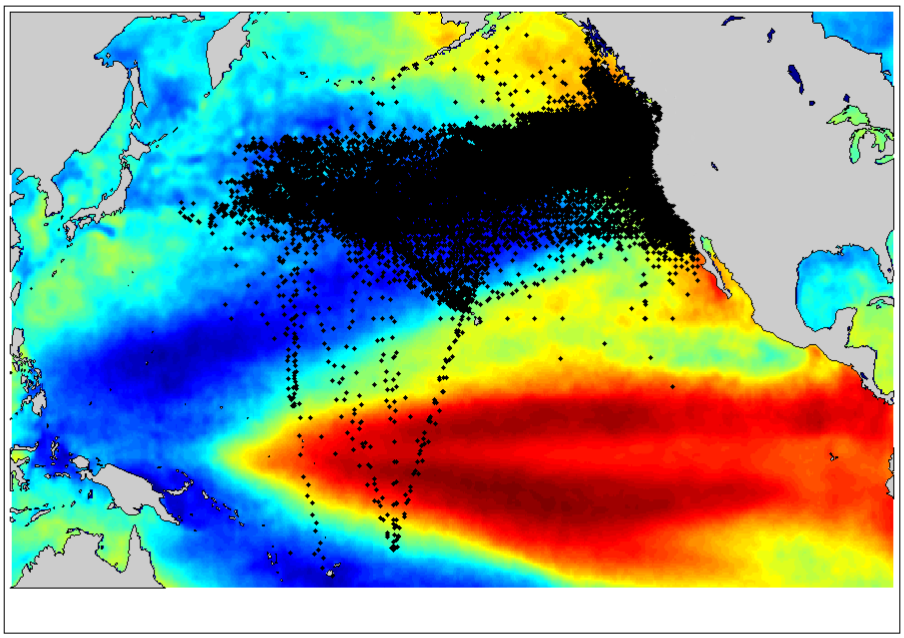

Notes: The heat map shows correlation between the one month lag of the Niño 3.4 index and sea surface temperature for each quarter degree latitude-longitude grid cell, as in Figure A5. Each point shows a daily observation of either fishing or transiting for a subset of the data from 1981 to 2010.

Table A7: Association of ENSO and Albacore Prices

	ln(albacore price)	ln(albacore price)
Niño 3.4	-0.015	-0.026
	(0.039)	(0.041)
L.ln(albacore price)	1.01***	1.01***
	(0.099)	(0.095)
Lag of Niño 3.4		0.053**
		(0.025)
Observations	36	36

Notes: The table shows results from estimating Newey-West regressions on annual time series data. The dependent variable is the log of the wholesale albacore price. In parentheses are Newey-West standard errors with 2 lags for autocorrelation. Significance indicated by: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A8: Association of ENSO and Fuel Prices

	ln(fuel price)
Niño 3.4 (t-1)	0.0032
	(0.0057)
Niño 3.4 (t-2)	-0.0056
	(0.012)
Niño 3.4 (t-3)	-0.0028
	(0.0089)
L.ln(fuel price)	0.99***
	(0.0096)
Observations	346

Notes: The table shows results from estimating Newey-West regressions on monthly time series data. The dependent variable is the log of the monthly average fuel price (marine diesel). In parentheses are Newey-West standard errors with 24 lags 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)
$\widehat{\text{Niño 3.4}}$	-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 (10) 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)
$\widehat{\text{Niño 3.4}}$	-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 (10) 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 (10) 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 (10) 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.