Expectations and adaptation to environmental risks

Jeffrey Shrader*

August 30, 2017

Abstract

Climate change is expected to have large, negative effects on the global economy. Adaptation by individuals and firms will determine, in part, how much damage ultimately occurs. This paper introduces a method for estimating forward-looking adaptation based on changes in expectations about the weather, provides conditions under which public forecasts provide good measures for these expectations, and formalizes identification of *ex ante* adaptation using *ex post* observations. To apply the method, I build a novel dataset of El Niño/Southern Oscillation (ENSO) forecasts and estimate adaptation by North Pacific albacore harvesters to ENSO-driven climate variation. The results show that, in this setting, nearly all of the effect of climate variation can be controlled through adaptation. Detailed, firm-level data allows for exploration of mechanisms, showing that vessels primarily adapt by timing entry into the fishery.

*Email: jeffrey.shrader@nyu.edu; New York University, 70 Washington Square South, New York, NY 10012. 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 comments on drafts of the paper. Additionally, I would like to thank Roy Allen, Richard Carson, Julie Cullen, Gordon Dahl, Matthew Gibson, Joshua Graff Zivin, Solomon Hsiang, Mark Jacobsen, Ben Miller, Francis Moore, David Pierce, Mark Rosenzweig, and Hendrik Wolff as well as seminar participants at the AAEA annual meeting, Berkeley ARE, Columbia University Sustainable Development, the Occasional Workshop in Environmental Economics, UC San Diego, 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

Climate change is predicted to have substantial, negative impacts on the global economy. The ultimate amount of damage will depend on both public policy to reduce emissions of greenhouse gases and on actions taken by individuals and society to adapt to a changing climate. Despite the role that adaptation plays in determining climate change outcomes, little is known about the total adaptation potential of climate-exposed industries or the economy. Moreover, much of what is known comes from analysis of *ex post* adaptation to experienced weather rather than *ex ante* adjustments made in expectation of climate change. Forward-looking adaptation is especially important because it helps individuals avoid damages before they occur, and studying this type of adaptation provides insight into the role of beliefs in determining behavior in environmental contexts. Changes in expectations about the climate suggest that such behavior will be an increasingly large part of the response to climate change going forward.

Estimating adaptation is challenging. Many individual mechanisms—such as choosing different inputs or altering consumption—might help reduce damage from a changing environment. An extensive literature has shown that individuals and firms do adapt to environmental changes along a number of dimensions.[1] The policy-relevant parameters, however, are the damage that results from changes in the environment net of *all* adaptation mechanisms and the aggregate cost of adaptation. Identification of these quantities 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 identify the overall effect of adaptation without reference to the underlying mechanisms. Following the seminal work of Dell et al. (2009), a recent literature has used average weather to estimate environmental effects *gross* of adaptation and used high frequency variation in weather to measure effects *net* of adaptation. Comparison of these estimates provides a measure of overall adaptation.[2] Surprisingly, given the evidence on individual adaptation mechanisms, these studies have generally found that total adaptation has little to no effect on output losses from

[1] For some recent examples, see Greenstone and Gallagher (2008), Neidell (2009), Graff Zivin and Neidell (2014), Graff Zivin et al. (2011), Deschênes and Greenstone (2011), Taraz (2015), and Barreca et al. (2016).

[2] These papers generally fall into one of two groups: those using short-run variation in the weather to get net-of-adaptation estimates and cross-sectional average weather to get gross-of-adaptation estimates, as in Dell et al. (2009, 2012); Hsiang and Narita (2012); Butler and Huybers (2013); Schlenker et al. (2013); Moore and Lobell (2014), and an approach that compares short-run variation to sub-sample average weather as in Burke and Emerick (2016). For a review, see Dell et al. (2014).

weather.[3]

In this paper, I use variation in individual expectations induced by public forecasts to identify total, *ex ante* adaptation.[4] This method contrasts with previous studies of total adaptation both in terms of the object of study and in the assumptions necessary for identification. Forecasts identify forward-looking adaptation rather than actions taken after an event occurs. Adaptation that occurs in advance of a change in the environment could be particularly important in many environmental contexts including climate change where disaster can result from a failure to avoid the bad state. Some researchers have questioned whether individuals will perform substantial *ex ante* adaptation in real-world settings (Mendelsohn, 2000). The method presented here allows for quantification of the degree of *ex ante* adaptation, and the empirical results show that such adaptation is practically important.

Intuitively, identification comes from an assumption that expectations about the weather only affect firm profit through input decisions—that there is no direct effect of information. Conditional on realizations of weather, then, forecasts contain only information available to firms before an event occurs, so the change in revenue with respect to a change in this information identifies the overall benefit of *ex ante* decisions. Under an additional assumption that the firms set all inputs before the state realizes, forward-looking adaptation is equal to total adaptation, and the method also identifies the direct effect of weather via weather realizations conditional on forecasts. Under these assumptions, the two estimates provide a complete picture of the damages a firm experiences due to weather.

The method shares the benefit of the work following Dell et al. (2009) that the researcher need not know the full suite of adaptation mechanisms available to an agent. In practice, this is because the estimation strategy regresses firm revenue on a forecast of a weather process and realizations of that process, and the forecast captures the "reduced form" or aggregate effect that forward-looking-input changes have on firm revenue. The method also has some unique benefits. First, by allowing the researcher

[3] Dell et al. (2009) find evidence for substantial adaptation in the gross domestic product–temperature relationship when comparing rich countries to poor countries. Over the last 50 years, however, Dell et al. (2012) shows that temperature effects on GDP have not weakened within income groups, a point reinforced by Burke et al. (2015).

[4] A small but growing literature in environmental economics is using forecasts to study forward-looking behavior. Neidell (2009) looks at the effect of pollution forecasts and public announcements on consumer behavior, Rosenzweig and Udry (2013) use monsoon forecasts to study optimal weather insurance for farmers, and Severen et al. (2016) ask whether farm land values have incorporated information from long-run climate forecasts.

to use firm revenue as the dependent variable, data requirements are reduced relative to envelope theorem-based methods that require profit (Hsiang, 2016). Second, the method allows for straightforward generalization to cases with discrete adaptation mechanisms, with the intuition again being that the reduced form effect averages over both continuous and discrete inputs. Third, by using a time varying measure of expectations, this strategy allows for empirical methods that alleviate omitted variable bias concerns. For instance, fixed characteristics of individuals or locations can readily be controlled for.

Applying the method, I estimate the degree of forward-looking adaptation to El Niño/Southern Oscillation (ENSO) by albacore tuna harvesters in the North Pacific. The empirical setting is particularly suitable for using forecasts to estimate adaptation. ENSO, a major source of global climate variation stemming from periodic but stochastic warming and cooling of the equatorial Pacific Ocean, was thought 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 months in advance of adverse events. Concurrent with these developments, the National Oceanic and Atmospheric Administration (NOAA) began a program to disseminate these forecasts to ENSO-exposed fisheries. The albacore fishery, historically a setting where output and profit declined substantially during ENSO, 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 were plausibly the only source of ENSO information available to albacore harvesters over the sample period.

Estimates show that the information in the forecasts is important to the fishery. The forecast has more than four times as large of an effect on revenue as does the realization of ENSO. Interpreting this through the lens of the model, the estimates suggest that forward-looking adaptation is large and effective in this setting. Harvesters are able to reduce the direct effect of ENSO to nearly zero, almost eliminating observable profit losses from this event. The results also show that if adaptation were ignored, estimates of the effect of ENSO on the fishery would be biased in two ways. First, the direct effect of ENSO on output and profit would be overstated because correlation between beliefs and outcomes causes some of the adaptation effect to be attributed to the direct effect. Overstatement of the direct damage from an environmental process when adaptation is ignored is a central concern when setting appropriate mitigation policy Mendelsohn et al. (1994). Second, the total effect of

ENSO would be understated because realizations of ENSO do not capture the adaptation effect that is only operating through expectations. This understatement, in a case where adaptation is costly, could lead to smaller than optimal policy responses.

Exploiting the richness of the spatially explicit, high-frequency, firm-level data, secondary results examine mechanisms by which the vessels use the forecasts to adapt. Overall, vessels respond to the forecasts by reducing their fishing effort during adverse periods. On the intensive margin, in anticipation of changes in ENSO vessels move closer to areas where albacore are expected to congregate given the weather change. This behavior suggests that ENSO mainly affects the fishery by increasing the uncertainty about where optimal fishing grounds will be located.

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, vessels choose to actively participate in the fishery much less often if ENSO is forecasted to be extreme. In contrast, the effect of realized ENSO conditional on the forecasts causes little or no change in behavior. Overall, the mechanism analysis supports the primary result. Revenue falls when the forecast of ENSO is high, but the behaviors engaged in by the firm are generally cost-saving measures, so 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 (2013) to determine whether the firms in this setting are risk averse. Intuitively, 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, since the past accuracy of ENSO forecasts (as measured by recent, historical mean 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.

ENSO is an important, global driver of medium-term climate that, in addition to fisheries, also affects health, civil conflict, agricultural productivity, worldwide commodity markets, and many other outcomes (Kovats et al., 2003; Hsiang et al., 2011; Solow et al., 1998; Brunner, 2002). The results from this paper show that economic agents can manage their risk from this climate process by making *ex ante* adaptation decisions. In the context of broader, global climate change, if vessels

are able to adapt to changing ocean temperatures due to climate change in a way that is similar to how they have adapted to ENSO, then the the results suggest that realized climate change damages might be greatly reduced. Caution should be exercised, however, since adaptation dynamics will certainly play an important role when extrapolating from the medium-term, cyclical variation considered in this paper to the longer-term changes caused by global climate change. Moving beyond the particular setting, the empirical method from this paper can be use to estimate adaptation in a number of industries to better inform impacts from ENSO and other weather phenomena. The novel dataset of ENSO forecasts created for the project can be used to assess adaptation to this climate process across the globe, and use of routine weather forecasts can help understand the scope for weather adaptation more generally.

Outside the context of environmental adaptation, the method discussed here also illustrates the contribution that analysis of forecasts of environmental processes can make to understanding long-standing problems in firm and consumer theory. 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 will allow for investigation of firm trade-offs in stochastic settings. Forecasts of environmental processes are well suited to study these issues not only because they are routinely used by firms and are easily observable by the researcher, but also because the processes about which the forecasts are being made are generally exogenous. This feature contrasts with other settings like finance where forecasts have the potential to endogenously change the state, complicating empirical analyses. Studying forward-looking behavior will likely become even more important in the future. Going forward, 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 such behavior.

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

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.[5] Formalizing this notion of adaptation helps one understand how to estimate both adaptation and total environmental impacts. In particular, a formal definition of adaptation will generalize from the single adaptation strategies or mechanisms that much of the economics literature has focused on—staying indoors on hot or polluted days (Neidell, 2009; Graff Zivin and Neidell, 2009), changing the mix of crops or the use of agricultural inputs (Rosenzweig and Udry, 2013; Hornbeck and Keskin, 2014), air conditioning (Barreca et al., 2016), or migrating (Deschênes and Moretti, 2009)—to the overall effect of adaptation on agent welfare.

The total effect of adaptation incorporates the effects of all adaptation mechanisms and identifying it is necessary for decomposing impacts into the effect that an agent chooses to control—the adaptation effect—and the residual portion that the agent chooses not to adapt away—the direct effect. This decomposition is important for understanding optimal public policy. If the scope for adaptation is small, then mitigation can have large, first order effects on the outcomes of agents. On the other hand, if adaptation is done in response to a pollutant, then even if adaptation potential is high, the costs of adapting should enter into the calculation of the pollution externality.

In this study, I will use expectations of agents to estimate the value of *total, forward looking adaptation*. This is the benefit to the firm of all behavioral responses that occur in advance of a change in the future state of the environment. Expectations drive such changes, as a consideration of the link between the adaptation mechanisms listed above makes clear. In making investment decisions or decisions like migration that involve high fixed costs, it is natural to characterize behavior as stemming from

[5]For examples of such a definition, see the Environmental Protection Agency's climate change website (`www3.epa.gov/climatechange/adaptation/`) or IPCC (2014). This study will primarily focus on individual consumer or firm adaptation.

an expectation that conditions will warrant the investment in the future. For behaviors that take time to set up or realize, expectations also play an obvious role. Even for short-run behavior, however, expectations are still important. This link is drawn explicitly by Neidell (2009). In the setting of that paper, public warnings are issued each day if pollution levels are forecasted to surpass a threshold. These forecasts are shown to have effects on how people choose their outdoor activities that day, highlighting the importance of expectations to even near-term decisions.

Formalizing this notion in a standard model makes the centrality of beliefs to total, forward-looking adaptation explicit and will lay the framework for econometric identification results. Consider a firm producing a univariate output at time t which is a function of weather as well as inputs that are chosen by the firm manager. Assume that the firm's production function is multiplicatively separable in terms of weather and inputs, so that at the beginning of each period, the firm's problem is to maximize expected profit[6]

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

Output price are denoted by by p, $\mathbf{c}$ is the J element 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.[7] Further assume that $f(\mathbf{x})$ is twice continuously differentiable and concave.[8] As is standard, a subscript on an expectation operator denotes the information set on which the expectation is conditioned, so $\mathbb{E}_{t-1}[g(Z_t)]$ is the expected weather this period conditional on information about the weather in all time periods up to and including period $t-1$. To emphasize the uncertain effect of weather on the production process, assume that the firm must choose each x_{jt} before the weather in period t is realized and that all x_j's are non-separable from Z.[9] Denote realized revenue by $y_t = pf(\mathbf{x}_t)g(z_t)$ and *ex ante* revenue as the expectation of this term

[6]Multiplicative separability is not a necessary assumption, but it improves the clarity of presentation and simplifies the estimating equation. For an extension of the model to non-separable weather, see Section A. I test the separability assumption empirically in Section 5.

[7]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 simply be replaced by Jacobian matrices.

[8]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 since the firm is not directly choosing Z.

[9]Additively separable inputs would not change in response to expected weather and are therefore not adaptations under my definition. For the more general model considering inputs chosen after weather has realized, see Section A.

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).[10]

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 nothing more or less than the set of changes in all inputs that are non-separable from weather. Optimized inputs 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

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

where, because this is a one-period problem, the time subscript on $\mathbf{A}$ has been dropped.

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 nominal return 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 total adaptation, 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

[10] In the empirical setting, the assumption of prices being uncorrelated with weather is testable and appears to hold. See Section 5.2.

researcher to be agnostic about the set of available mechanisms since expectations will fully capture the reduced form effect of all forward-looking adaptation. The downside to using expectations is that one cannot analyze the contribution of each adaptation mechanism to the overall level of adaptation; one would need an instrument for each input in order to do this. Relatedly, the first order conditions suggest that adaptation could be inferred from reductions in the direct effect of weather on profit over time. If adaptation potential increases, for instance, due to increasing productivity or decreased costs, then the direct effect of weather on profit should decrease. This is the empirical strategy pursued by Hornbeck and Keskin (2014).

Third, the continuity assumption is not necessary for the definition of adaptation. For discrete adaptations like technology adoption or changes in land use, the derivatives in Equation (3) can be replaced by differences. In this case, adaptation is the change in inputs, broadly defined, in response to changes in the environment. Handling the case of discrete inputs is an important feature of any empirical method for studying adaptation in light of the dramatically different welfare implications of the continuous versus discrete cases. Continuous adaptation is, in classical models, welfare neutral (a direct result of the envelope theorem) while discrete adaptations are potentially welfare improving as shown by Guo and Costello (2013). Estimates of the value of adaptation using expectations and revenue are robust to discrete inputs, as will be discussed below.

Continuing the formalization within the continuous model, the value or benefit of adaptation is the adaptation vector multiplied by the revenue value of those changes, denoted

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

where arguments of the maximized output and choice variables have been suppressed for clarity. Estimating this value is the primary goal of this study. Such an estimate is important for many reasons. Under the assumption of continuous adaptation, the value of adaptation provides information on adaptation costs, it provides information on how much adaptation contributes to revenue for the firm, and it is crucial, in general, for estimating the total effect of weather on the firm.

Also important for policy is the direct effect of weather. In the context of the

model, the direct effect of weather is

$$\frac{\partial \mathbb{E}_{t-1}[y_t^*]}{\partial \mathbb{E}_{t-1}[g(Z_t)]}. \tag{5}$$

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.

From the model, one can see that if a researcher observes the expectations of agents and has access to *ex ante* data, then both the value of adaptation and the direct effect of weather can be estimated. In general, neither of these conditions is likely to hold. The next two sections show that identification can still be achieved with *ex post* data and a well-chosen proxy for agent beliefs.

2.2 Identifying *ex ante* adaptation with observed data

This section formalizes identification of the value of adaptation and the direct effect of weather using *ex post* observable data. It is assumed that the researcher has access to accurate measures of agent expectations about the weather. This assumption is relaxed in the next section. Here, I show parametric identification results with a known functional form for the function of weather, g, and I assume that weather is multiplicatively separable from inputs. For the more general case with non-separable inputs and non-parametric identification, see Section A.1.

Intuitively, identification is driven by the assumption that, conditional on expectations, realized weather does not influence the input decisions made by firms at the beginning of each period. Under this assumption, holding expectations fixed also holds inputs (adaptation) fixed. Varying the realization of weather in this case traces out the direct effect of weather on revenue. Changes in expectations holding realizations fixed have a complementary effect. Only forward-looking inputs are varied in this case, identifying the output effect of adaptation.

More formally, inputs are a function of expected weather and not realized weather, so $\mathbb{E}_{t-1}[f(\mathbf{x}^*)] = f(\mathbf{x}^*)$. Thus, the direct effect is identified exactly by *ex post* data because $\partial y_t / \partial g(z_t) = pf(\mathbf{x}^*) = \partial \mathbb{E}_{t-1}[y_t] / \partial \mathbb{E}_{t-1}[g(Z_t)]$.

For identification of the adaptation effect, note first that with respect to the information at time $t-1$, $\partial \mathbf{x}^* / \partial \mathbb{E}_{t-1}[g(Z_t)]$ 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 partial derivative of *ex ante* output with respect to expected weather. For estimation, a regression of revenue on $g(z_t)$ and $E_{t-1}[g(Z_t)]$ will return unconditional averages of these derivatives. These averages identify the derivatives of interest after an application of the law of iterated expectations.

2.3 Using public forecasts to measure beliefs

Given the identification argument presented above, the ideal estimating equation to measure adaptation and direct effects from 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.

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 expectations of private agents in a way that is amenable to estimation.

Professional forecasts will provide a good measure of agent beliefs under the assumptions that the forecasts are public, that agents are maximizing expected profit, and to the degree to which the forecasts capture the full information available to agents. Under these conditions, it can be shown that forecasts are good proxies for agent expectations.

To see this, denote the public forecast as $\widehat{g(z)}$, and consider the public forecast as a proxy for the private expectation (Wooldridge, 2010, ch.4). The first condition for a good proxy is that it is redundant with the variable being proxied for. In this case, 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 $\mathbb{E}\left[y|g(z), \mathbb{E}^p[g(Z)], \widehat{g(z)}\right] = \mathbb{E}\left[y|g(z), \mathbb{E}^p[g(Z)]\right]$. 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 forecast 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). Writing public forecast as a linear projection of private beliefs

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

this condition can be formalized as saying that if the researcher estimates

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

then the covariance between realized weather and the error term from Equation (7) needs to be zero. In other words, one needs $\mathbb{E}[g(z_t)\xi_t] = 0$, assuming that exogeneity holds for the true Equation (6). 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. Elaboration on this condition can be found in Section A.5.

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.

An alternative approach to measuring agent expectations 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.[11] In cases where the relevant

[11]The error in this approximation can be bad in extreme cases. For instance, if agents have

stochastic 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 the 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. Finally, average weather cannot be used in cases where the relevant climate shifts are measured in terms of anomalies (as in the empirical setting of this study). 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.

2.3.1 Violations of forecast proxy conditions

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

Maintaining the assumptions that forecasts are public and that agents are fully sophisticated but making no assumption about the relationship between the public and private forecasts, 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 $\operatorname{plim} |\tilde{\alpha}_1| = \left|\alpha_1 + \alpha_2 \frac{\mathbb{C}\text{ov}(\xi, g(z))}{\mathbb{V}(g(z))}\right| > |\alpha_1|$. The magnitude of the coefficient is larger because the sign of α_1 should be the same as the sign of α_2 and because of the positive covariance between ξ and $g(z)$. Therefore, the direct effect will be over-estimated, leading to downward bias on the relative degree of adaptation.

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

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.

the sign of the bias of the estimator employed by the forecaster or agent.

If the firm and forecaster information sets are partly disjoint or if the firm creates its own forecasts but with a smaller information set than the public forecaster, 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, using the public forecast for estimation will return the correct sign on the adaptation effect and will help reduce the omitted variable bias from ignoring adaptation.

3 Background and data

3.1 Albacore fishing, ENSO, and ENSO forecasting

Three attributes of the North Pacific albacore fishery make it an ideal setting to study adaptation. First, ENSO has a substantial effect on the fishery both because ENSO causes substantial changes to the weather and oceanic conditions of the North Pacific and because albacore are sensitive to those changes. Second, NOAA issues forecasts directly to albacore harvesters in the fishery, and interviews with harvesters indicate that these forecasts are utilized. 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). Also, the U.S. harvesters studied here account for a small part of the global albacore tuna output, mitigating concerns about aggregate output price effects from ENSO, and the primary variable cost comes from diesel fuel, a globally traded and produced commodity.

Albacore (*Thunnus alalunga*) typically follow oceanic fronts with strong temperature gradients and stay in waters with sea surface temperature between 15 and 20°C (Childers et al., 2011). The temperature preferences of albacore make them highly responsive to changes in climate. The preferences of the albacore have led harvesters to develop rules of thumb based on sea surface temperature ranges when determining where to try to catch fish (Clemens, 1961; Laurs et al., 1977). Since the mid-1980s, scientists and harvesters have become increasingly aware of the influence of other factors in determining albacore location, including water color and clarity, but temperature remains an important choice variable for harvesters when determining fishing location (Laurs et al., 1984; Childers et al., 2011).

ENSO affects the temperature of the North Pacific (see Figure 6) and oceanic structures like temperature gradients. These shifts make it harder for vessels to locate albacore (Fiedler and Bernard, 1987).[12] 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 in order to pursue crabs and other pelagic species less affected by ENSO conditions (Wise, 2011; McGowan et al., 1998).

The average fishing trip is about two weeks long, and 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).[13]

Skillful forecasts of ENSO were developed starting in the mid 1980s. An early ENSO forecast based only on atmospheric modeling was published by Inoue and O'Brien (1984). Cane et al. (1986), a group of researchers at the Lamont-Doherty Earth Observatory (LDEO), published the first coupled ocean-atmosphere forecast, termed LDEO1. In the late 1980s, NOAA's Climate Prediction Center (CPC) began to produce a statistical forecast of ENSO based on Canonical Correlation Analysis (CCA). A stated goal of the LDEO forecasting group was to produce unbiased forecasts of ENSO (Chen et al., 2000).

Starting in June 1989, the LDEO forecast was issued publicly in NOAA's Climate Diagnostics Bulletin, a publication of global climate information and medium term climate forecasts. The Climate Diagnostics Bulletin incorporated additional ENSO

[12] 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 13 and rule it out as an explanation of the short-run results.

[13] 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." http://www.pmel.noaa.gov/tao/proj_over/taohis.html

forecasts as they were published, starting with the CCA forecast in November 1989.[14] Today, the Bulletin publishes 21 ENSO forecasts on a monthly basis. See Appendix B.1 for more information on the content of the Bulletins. Analyses of forecast accuracy and performance over time can be found in Barnston et al. (2010, 2012).

At nearly 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.[15]

For this paper, I focus on the effects of the 3-month-ahead ENSO forecast. The use of this forecast is primarily due to data constraints—it is the only forecasting horizon that I observe over the full sample period—but it is also because of practical considerations. The Bulletin forecasts are typically released a month after they have been generated, so a three month ahead forecast is, practically, a one or two month ahead forecast from the perspective of the fisher. Given the timing of ENSO effects being felt in the North Pacific and typical trip length, this forecast horizon is likely to be the relevant one for fishing decisions.

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.

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. Anomalies are calculated with respect the thirty-year average temperature. This study uses the 1971-2000 average. Following Trenberth (1997) and NOAA, I classify El Niño and La

[14] For examples of these historical Bulletins, one can see the archive going back to 1999 at the following link: `http://www.cpc.ncep.noaa.gov/products/CDB/CDB_Archive_html/CDB_archive.shtml`

[15] 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. See `http://www.seaviewfishing.com/News.html`

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

Panel A: Pre-forecast sample (1981-May 1989)

	Mean	St. Dev.	Obs.
Catch per month (fish)	163.30	611.15	26,415
Catch weight (pounds)	1,079.31	5,849.68	26,415
Niño 3.4 index	0.01	1.02	26,415
Vessel length (ft)	50.50	9.63	26,385
Diesel price (2001 $)	1.95	0.68	21,710
Albacore price (2001 $)	1.35	0.27	20,061

Panel B: Post-forecast sample (June 1989-2010)

	Mean	St. Dev.	Obs.
Catch per month (fish)	264.92	980.86	69,057
Catch weight (pounds)	3,081.94	12,687.92	69,057
Niño 3.4 index	0.16	0.81	69,057
3 month-ahead Niño 3.4 forecast	0.09	0.58	69,057
Vessel length (ft)	55.01	18.73	66,444
Diesel price (2001 $)	1.72	0.79	67,483
Albacore price (2001 $)	1.08	0.23	62,894

Notes: Averages, standard deviations and number of observations for primary variables in the dataset are shown for the pre-forecast (panel A) and the post-forecast (panel B) samples. Between 1981 and 2010, the dataset contains 2,125 unique vessels.

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 generally 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 2002 to the present. 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 from 1981 to 2010. All fishing days are observed, with additional information provided for some transiting and port days (these latter data do not appear to be consistently reported). For each fishing day, the logbooks report

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, and the number of troll lines used. At the trip-level, the logbooks report vessel length, departure and arrival port, and total weight of catch for the trip. Landing port is matched to the Pacific Fisheries Information Network (PacFIN) database on annual albacore sale prices for 1981 to 2010. 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).

The vessels in the sample use #2 marine diesel fuel. Where available, the price for this fuel is used for cost calculation, but the price for this exact fuel type is not available over the full sample. From 1983 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.[16] All prices are pre-tax if possible. See Appendix B.3 for further details. All prices have been deflated to 2001 dollars using the monthly core consumer price index from the U.S. Bureau of Labor Statistics available from the Federal Reserve Bank of St. Louis' FRED database.

Finally, full costs, expenditures, and revenues for a panel of 35 albacore harvesters were 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

To estimate the effect of ENSO on the fishery one would ideally regress output on the forecast and realization of ENSO, both transformed by a known function g, as

[16] Available online from `www.psmfc.org/efin/data/fuel.html`.

in Equation (6). Since the function g is unknown in this case, I will first present non-parametric results, and I will then give regression results with a theoretically motivated parametric specification.

For the latter case, assuming that vessels are well adapted to "typical" climate conditions suggests that profit should be highest when ENSO anomalies are neither high nor low—in other words, when neither a La Niña nor an El Niño is occurring. In that case, unexpected deviations in either direction will cause loss in profit relative to the zero-anomaly case, so the relationship between the ENSO, as measured by the Niño 3.4 index, and profit or revenue should be concave (recall that the model shows that the direct effect is the same for both profit and revenue). A simplifying assumption is that this relationship is symmetric for positive or negative ENSO events. This theoretical relationship suggests that a quadratic function for g is appropriate. The lag between changes in ENSO in the equatorial Pacific and the effects being felt in the North Pacific suggests that this function should be in terms of the lag of ENSO. Putting this together, let

$$g(z_{t-1}) = \gamma + \beta_1 z_{t-1} - \beta_2 z_{t-1}^2, \tag{8}$$

where γ is some positive constant sufficiently large to ensure that vessels would like to enter the fishery and z is the Niño 3.4 index. Because the Niño 3.4 index is centered around zero, the assumption that vessels are well adapted to normal conditions implies that $\beta_1 = 0$, so a simplified equation could exclude this term.

Given this function of weather, if agents are forming distributional beliefs about ENSO, then the correct forecast term to include would be $\widehat{g(z_{t-1})} = \gamma + \beta_1 \mathbb{E}_{t-h}[Z_{t-1}] - \beta_2 \mathbb{E}_{t-1}[Z_{t-1}^2]$, 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

$$\widehat{g(z_{t-1})} = \gamma + \beta_1 \mathbb{E}_{t-h}[Z_{t-1}] - \beta_2 \mathbb{E}_{t-h}[Z_{t-1}]^2 \tag{9}$$

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 dis-

tribution. Then

$$\mathbb{E}[g(Z)] = \gamma + \beta_1 \mathbb{E}_{t-h}[Z_{t-1}] - \beta_2 \mathbb{E}_{t-h}[Z_{t-1}^2] \tag{10}$$

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

$$\mathbb{E}[g(Z)] - g(\mathbb{E}[Z]) = \beta_2(\mathbb{E}_{t-h}[Z_{t-1}]^2 - \mathbb{E}_{t-h}[Z_{t-1}^2]) = \beta_2 \mathbb{V}_{t-h}(Z_t) \tag{11}$$

If one assumes that Z_t has constant variance over time, then (11) is constant so the difference between the two measures will be absorbed by the intercept term. Then, 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 Appendix C Figure 5, 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 relative to the movement in the mean. Much of the research on climate change has focused on uniform shifts in the location of the weather distribution, but climate change is expected to have effects on higher moments of weather as well. Therefore, future work would benefit from using distributional forecasts to assess adaptation to changes in the full distribution of weather.

Putting all elements together, the full estimating equation is

$$y_{it} = \beta_0 + \beta_1 z_{t-1} + \beta_2 z_{t-1}^2 + \beta_3 \hat{z}_{t-1} + \beta_4 \hat{z}_{t-1}^2 + \mathbf{x}_{it}'\alpha + \varepsilon_{it} \tag{12}$$

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 indicated by the slope of the $\hat{z}$ terms relative to that of the z terms. This will be considered formally in Section 5.3, but intuitively, the higher the magnitude of β_4 relative to β_2, the greater the adaptation.

5 Results for ENSO effects and adaptation

5.1 Adaptation, direct effect, and total effect of ENSO

The timing of the release of public ENSO forecasts in 1989 allows for an initial assessment of adaptation by comparing the effect of ENSO before forecasts were released

to the effect after the release. Under the assumption that ENSO was unforecastable, agent expectations in this period would be climatological or unchanging over time. In that case, the effect of ENSO on output captures the effect absent any forward-looking adaptation. After 1989 and the release of forecasts, the relationship between ENSO and output should capture an average of the direct effect and the forward-looking adaptation effect. In this case, one would expect the relationship to be attenuated relative to the pre-forecast period if adaptation is occurring.[17]

Figure 1 gives results from implementing this method. The figure shows local linear regressions between output (the y-axis) and the one-month lag of the Niño 3.4 index (x-axis) for the period before forecasts were released (1981-May 1989) in red, and the period after forecasts were released (June 1989 to 2010) in blue. Both the output and Niño 3.4 index measures are residuals from regressions on month indicators to remove seasonality.

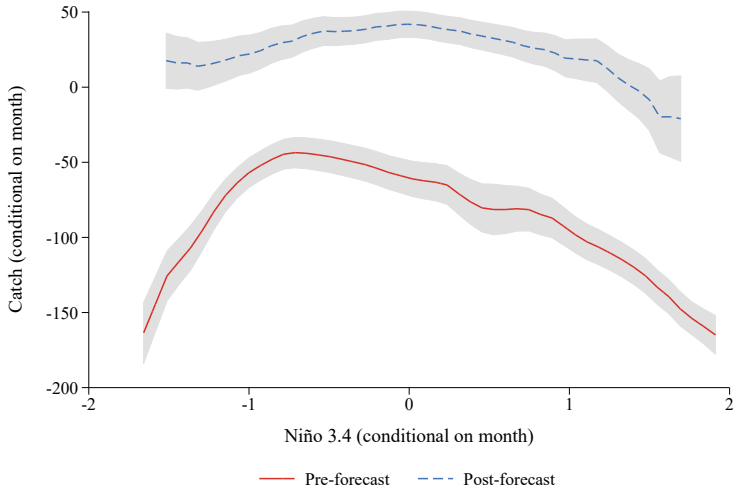
Before the introduction of forecasts, harvesters experienced large declines in catch at both high and low levels of ENSO. Average catch in a month during this period was 155 fish, so going from "normal" conditions (index value of 0) to a moderate El Niño (index value of 1) was associated with a decrease in catch of about a third. The losses were even steeper for extreme negative values of the index (La Niña events). This result shows that ENSO was an important driver, historically, of catch in the fishery.

In the period after forecasts were released, the relationship between ENSO and catch flattens substantially and the effect becomes more symmetric about zero. Overall, catch per month has risen in the fishery between the 1980s and the present for many reasons. Identification of the adaptation effect comes not from this level shift in catch, however, but from the change in curvature between the solid and dashed lines. The reduced curvature after forecasts were released provides initial evidence that adaptation to ENSO is occurring in the fishery.

This figures does not, however, give a complete measure of adaptation. The relationship after the release of the forecasts is a combination of the direct effect of ENSO and the effect of adaptation by the firm. Because realizations of ENSO are not perfectly correlated with forecasts, this combination will, in general, be attenuated

[17]The assumption of unforecastability is likely too strong, even in light of evidence presented in Section 3.1 that in the 1980s ENSO was not consistently observed, much less predicted. ENSO anomalies exhibit autocorrelation, so once an ENSO event begins, it is likely that it will last for the rest of the year. Therefore, this evidence should be considered a lower bound on the effect from adaptation.

Figure 1: Output and ENSO before and after forecasts



Notes: Each line shows a local linear regression (Epanechnikov kernel with bandwidth of 0.38) of catch on the Niño 3.4 index the previous month. Both variables are residualized on month of year to remove seasonality. The red, solid line uses the sample from 1981 to May 1989 before ENSO forecasts were released. The blue, dashed line uses the sample from after forecasts were released in June 1989 until 2010. Shaded areas give the 95% confidence intervals.

relative to the true total effect. The formal estimation strategy isolates the direct effect from ENSO by regressing changes in the Niño 3.4 index on catch, controlling for expectations, and it isolates the forward-looking adaptation response using forecast changes holding realizations fixed. The total response by the firm to ENSO is the sum of these two effects. A more careful analysis of ENSO effects in a regression framework can perform this decomposition while also including control variables for fixed vessel or time characteristics.

Table 2 gives results from implementing the formal identification strategy. Each column shows estimates of versions of Equation (12) using monthly data. The dependent variable in the first two columns is the number of fish caught per month by each vessel, in the third column it is the log of the number of fish caught, and in the fourth column it is revenue. 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, using a uniform kernel, a distance cutoff of 30km, and 2 year lags for autocorrelation (Conley, 1999).

Table 2: Effect of ENSO on catch and revenue

	(1) Catch	(2) Catch if fishing	(3) Revenue	(4) Catch
Niño3.4$_{t-1}$	103.7***	-3.69	1153.4**	29.8
	(36.6)	(91.4)	(473.0)	(22.2)
Niño3.4$^2_{t-1}$	-16.3	-107.6***	-66.4	-32.1***
	(15.7)	(40.2)	(200.9)	(11.5)
$\widehat{\text{Niño3.4}}_{t-1}$	-97.5***	-132.3	-1435.4***	
	(31.3)	(86.3)	(373.4)	
$\widehat{\text{Niño3.4}}^2_{t-1}$	-72.5***	95.6*	-745.3**	
	(27.4)	(53.2)	(333.9)	
Vessel FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Observations	69,057	12,430	62,894	69,057
R^2	0.079	0.16	0.065	0.077

Notes: The table shows results from estimating equation (12) on monthly data. The dependent variable in each model is indicated at the top of the column. Catch is the total number of fish caught per month by a vessel and revenue is the total ex-vessel value of that catch. Catch if fishing is the sub-sample of observations when vessels are active in the fishery and engaged in fishing in a given month. 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 2 year lags for autocorrelation. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

For all but the last column, four coefficients are reported, corresponding to β_1 through β_4 from Equation (12). The *Niño3.4* and *Niño3.4*^{2} coefficients give the effect on catch or revenue of a 1°C change in the the Niño 3.4 index. The $\widehat{\textit{Niño3.4}}$ and $\widehat{\textit{Niño3.4}}^2$ coefficients give the effect from a forecasted change in ENSO.

In the last column, only the Niño 3.4 index measure is included. This column shows the inference that would result from naïvely estimating the effect of ENSO on the fishery while ignoring expectations. The results indicate that ENSO has a moderate, negative effect on catch. A one standard deviation change in the Niño 3.4 index is about 1, so column 4 says that a typical change in the Niño 3.4 index leads

to a loss in catch of about 30 fish per month. Average catch is about 260 fish per month, so this represents a little more than 10% change in catch.

Without including forecasts, however, this result does not give a complete or accurate picture of the effects of ENSO in the fishery. Column 1 adds variables for the forecast of the Niño 3.4 index. One can see that predicted changes in ENSO actually have a much larger effect on output than realized changes. In particular, changes in information lead to a change in output more than four times larger than a comparable change in realized ENSO.

Summing the effects from both realized and forecasted ENSO, moving from normal conditions to a moderate El Niño (Niño 3.4 index of 1) leads to a 30% decline in output, on average, for a vessel. The effect from a change in ENSO conditional on the forecast, however, is reduced substantially. Comparing these results to column 4, the naïve method overstates the effect of a change in realized ENSO by a factor of 2. This illustrates the bias in climate damage estimates that can result from ignoring adaptation, as argued by Mendelsohn et al. (1994). In addition, the total effect is underestimated by a factor of 3. Since adaptation is, in general, costly, this high degree of adaptation also has bearing on welfare analysis from this process.

Column 2 looks at the effect of ENSO on catch conditional on a vessel choosing to fish in a given month. Vessels typically only choose to fish for albacore one-fifth of the months that they are in the fishery. One can see that conditional on choosing to go fishing, forecasts still have a substantial effect on catch—a 1 unit change in the forecast of ENSO causes about a 20% decline in the number of fish caught in this case—but the effect of a realized change in ENSO is much greater relative to the full sample results.

Column 3 shows estimates using revenue (in constant 2001 dollars) as the dependent variable. Revenue information is not available for the full dataset, either because the logbook record is missing information on the weight of the fish caught or because the vessel offloads fish at a port outside of California, Washington, or Oregon where albacore price is observed. The results reported in this table use imputed weight where weight is missing. The effect of this imputation is assessed in robustness Table 12. The missing values in revenue lead me to prefer the results using number of fish caught, but comparison between columns 1 and 3 shows that the results are qualitatively similar between the two samples. This result provides initial evidence that albacore prices are not changing in response to changes in ENSO, a topic that will be taken up in detail in Section 5.2.

Overall, these estimates provide evidence that beliefs correlated with the public ENSO forecasts are important for output and revenue in the fishery. Assessing these estimates in the context of adaptation requires the additional identifying assumptions laid out in Sections 2.2 and 2.3. Support for these assumptions is discussed in the following sections, and formal calculation of the adaptation effect is carried out in Section 5.3

5.2 Price effects and profit

Measuring adaptation with output and revenue, as is done in the previous section, is convenient from the standpoint of data availability and as the theory makes clear, it might also be necessary in cases where a substantial portion of the adaptation mechanisms are discrete. If profit is continuous in all adaptation mechanisms, then an application of the envelope theorem shows that the marginal profit value of adaptation is zero. In this case, estimates using profit as the dependent variable can return the direct effect of weather but not an explicit measure of adaptation. On the other hand, if some adaptation mechanisms are discrete, as in Lemoine and Traeger (2014), then the profit effects of adaptation will be greater than or equal to zero in an optimizing model (Guo and Costello, 2013). In general, using profit as the dependent variable in a regression with only weather on the right hand side will yield estimates that are an average of the direct effect and the effect of discrete adaptations.

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 number or wages of crew. The one input that can be consistently calculated is movement during fishing trips. Appendix Section B.5 has 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. To arrive at movement costs, I multiply movement by the real price of fuel, based on port-level records. 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.

Table 3 compares the effect of forecasted and realized ENSO on revenue and revenue net of movement costs, both for a consistent sub-sample where profit is ob-

Table 3: ENSO effects on partial profit

	(1) Revenue	(2) Net revenue
Niño3.4$_{t-1}$	1171.4**	1003.2***
	(473.2)	(373.7)
Niño3.4$^2_{t-1}$	-73.0	-85.4
	(200.8)	(156.0)
$\widehat{\text{Niño3.4}}_{t-1}$	-1439.3***	-1089.5***
	(373.6)	(308.6)
$\widehat{\text{Niño3.4}}^2_{t-1}$	-752.4**	-643.8**
	(333.5)	(262.4)
Baseline FE	Yes	Yes
Observations	62,868	62,868
R^2	0.066	0.042

Notes: The table shows results from estimation using monthly data. The dependent variable is monthly average profit. 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 2 year lags for autocorrelation. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

served. As predicted, the magnitude of the effect of forecasted changes in ENSO falls for partial profit. Theory suggests that since movement is an intensive adaptation mechanism, the profit effect should be zero for the anticipated component. The results support this conclusion, with the profit changes due to movement falling by about 15% for anticipated changes in ENSO. For unanticipated realizations, the linear term also falls by 15% but the square term increases in magnitude by a similar amount.

These changes in profit are coming primarily through changes in firm behavior rather than through changes in albacore or fuel prices. The lack of observable change in albacore price in response to changes in ENSO can be inferred from a comparison of the revenue and output results. Running a more explicit analysis of changes in ENSO on the average time series for albacore and fuel prices shows that ENSO is weakly, negatively associated with both prices. These results can be found in Table 14.

5.3 Quantifying the importance of adaptation

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,

$$V_n(\mathbf{A}) = \frac{V(\mathbf{A})}{d\mathbb{E}_{t-1}[y_t^*]/d\mathbb{E}_{t-1}[g(Z_t)]}. \tag{13}$$

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.

$$\frac{d\mathbb{E}_{t-1}[y_t^*]}{d\mathbb{E}_{t-1}[g(Z_t)]} = \frac{\partial\mathbb{E}_{t-1}[y_t^*]}{\partial\mathbf{x}_t^*} \cdot \frac{\partial\mathbf{x}_t^*}{\partial\mathbb{E}_{t-1}[g(Z_t)]} + \frac{\partial\mathbb{E}_{t-1}[y_t^*]}{\partial\mathbb{E}_{t-1}[g(Z_t)]} \tag{14}$$

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 value 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.

Estimating the normalized value of adaptation using the parametric specification in Equation (12) 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.[18]

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 value of adaptation. First, one can condition on being away from the point of zero slope when estimating the expectations in Equation (13). This method is convenient, but it also has interpretability. If the functional form of the relationship between the level of ENSO and adaptation is such that more extreme events are harder

[18]The value of $V_n(A)$ for all observations of Niño 3.4 can be found in Appendix Figure

Table 4: Quantifying the effect of adaptation

Estimator of $V_n(A)$	(1) Catch	(2) Catch if active	(3) Revenue
Average conditional on	0.86	0.40	0.81
$\lvert$Niño3.4$\rvert > 0.5$	(0.35,1.37)	(-0.05,0.85)	(0.10,1.52)
Limit as Niño3.4 $\to \infty$	0.82	0.47	0.92
	(0.47,1.17)	(0.12,0.82)	(0.43,1.41)
Median	1.01	0.78	1.02
	(0.48,1.54)	(0.39,1.25)	(0.47,1.57)

Notes: The table shows results from three estimators of Equation (13) using monthly data. The dependent variable in each column corresponds to a model from Table 2. 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.

(or, less plausibly, easier) to adapt to, then conditioning on progressively high values of ENSO will reflect that change. In practice, I condition on the Niño 3.4 index being greater than 0.5 of less than -0.5, the cut-off for declaring an El Niño or La Niña, respectively.

Second, one can calculate the median of $V_n(A)$ using the empirical distribution of ENSO. The median is less subject to outliers caused by division by zero, and even if the true distribution of ENSO is of the family with no first moment (for instance, the normal distribution), then the median still exists. 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 300 bootstrap replications are shown.

Finally, for the parametric specification used in the baseline results, one can take a limit of the numerator and denominator of Equation (13) as Niño 3.4 goes to infinity. Because of the parametric assumption used to estimate the baseline results, this limit is not a function of ENSO, and $V_n(A)$ simplifies to be $\beta_2/(\beta_2 + \beta_4)$, where the coefficients are those from Equation (12). This method has the advantage that standard errors can be easily calculated using the delta method, under the assumption that $\beta_2 \neq 0$. Given the quadratic estimating equation and the estimated parameters, the limit-based estimate of $V_n(A)$ will agree with the conditional average-based estimate for a sufficiently wide interval of excluded Niño 3.4 values.

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). For intensive-margin adaptation, the conditional average estimate is only marginally statistically different from zero, but for the other two estimators are highly significant. In none of the full adaptation cases can 100% adaptation be rejected at conventional significance levels.

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 2.2 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 buoy measurement, so the measurement error in the forecasts should lead to attenuation of the forecast coefficient.

5.4 Robustness

Three parametric assumptions underlying the estimates can be assessed. First, the quadratic functional form chosen for the estimating equation is tested nonparametrically in Figure 1. In both the pre and post-forecast samples, the overall effect of ENSO on output appears to be quadratic. Second, the constant variance assumption is tested by calculating a rolling variance of the Niño 3.4 index in Appendix Figure 5. Aside from a period of high variance in the late 1990s and early 2000s, this assumption appears to roughly hold. Re-estimation of the baseline specification excluding this period is done in Table 6, and the results are largely unchanged. Finally, the assumption of multiplicative separability is tested in Table 13, Column 5 by including an interaction between the forecast and realization of ENSO. High correlation between the interaction term and the square terms prevents separate identification of these effects. Note that this term cannot be used to assess forecast quality under the assumption that production is concave in ENSO. In this case, firm profit is highest if ENSO always turns out to be at whatever point corresponds to the peak of this

concave function. For instance, if the firm is well adapted to normal conditions, then profit should be highest at a Niño 3.4 near zero, regardless of whether the forecast is accurate or not.

Table 13 also provides checks of robustness to changes in controls and the method of standard error calculation. In Column 1, the separate vessel and year fixed effects are replaced by a set of vessel-year fixed effects. These more flexible controls do not appreciably change inference. Column 2 adds vessel-specific linear trends, again to negligible effect on inference. Trends could be important since catch is rising, on average, over time, and forecast quality is also changing over time (Appendix Figure 4). Another test to rule out trends as spuriously driving these results is reported in Figure 11, which replaces the level of ENSO with the difference in ENSO between the previous month and the month before that. Output has a concave and symmetric relationship with the change in ENSO.

Column 3 clusters standard errors at the year-month level. ENSO is a group shock, and forecasts are released each month, so this level of clustering more closely matches the level of aggregation of the exogenous shock. Inference is slightly less precise in this case—two variables go from being significant at the 1% level to being significant at the 5% 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 6), so year-month clustering is likely to be too conservative.

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. Controlling for a quadratic in 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. The conclusion of Lehodey et al. (2003) is strongly supported by the data, with year-ago ENSO values having a comparable effect on catch to the contemporaneous measures.

Table 6 contains two more variations in specification and two sample restrictions. The specification in Column 1 includes only the square Niño 3.4 terms. The theoretical motivation for the quadratic specification discussed in Section 4 suggested that excluding the linear terms could be appropriate. The significant linear terms in the baseline model show that this conclusion is likely untrue, but the results are qualitatively similar if the linear terms are forced to be zero. Note that the calculation of $V_n(A)$ is simplified in this case because the ratio will not be a function of ENSO in

Table 5: Robustness to clustering and controls

	(1) Catch	(2) Catch	(3) Catch	(4) Catch	(5) Catch
Niño3.4_{t-1}	107.5***	106.8***	103.7**	87.1**	109.8***
	(33.1)	(36.6)	(49.9)	(35.7)	(36.6)
Niño3.4^2_{t-1}	-18.1	-15.9	-16.3	-14.4	28.6
	(14.3)	(15.6)	(25.2)	(15.5)	(49.8)
$\widehat{\text{Niño3.4}}_{t-1}$	-92.9***	-94.6***	-97.5**	-115.6***	-101.9***
	(29.6)	(31.5)	(47.5)	(39.3)	(31.9)
$\widehat{\text{Niño3.4}}^2_{t-1}$	-74.1***	-76.9***	-72.5**	-80.4***	-21.9
	(24.5)	(27.8)	(35.1)	(29.9)	(36.8)
Vessel trend		-101.5***			
		(30.6)			
Niño3.4_{t-12}				61.5***	
				(21.1)	
Niño3.4^2_{t-12}				-68.3***	
				(13.1)	
$\text{Niño3.4}_{t-1} \times \widehat{\text{Niño3.4}}_{t-1}$					-98.4
					(93.6)
FEs	Vessel-year	Baseline	Baseline	Baseline	Baseline
SEs	Spatial	Spatial	Year-month cluster	Spatial	Spatial
Observations	69,057	69,057	69,057	69,057	69,057
R^2	0.10	0.079	0.10	0.081	0.079

Notes: The table shows results from estimating versions of equation (12) on monthly data. The dependent variable in each model is the monthly catch, where catch is the number of fish caught. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects unless otherwise noted. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 2 year lags for autocorrelation, unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

this case.

Column 2 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, if anything, strengthens the results. The sample restriction in Column 3 has already been discussed.

Columns 4 adds the one-month lag of catch. The baseline estimates use two year lags to account for autocorrelation in the residuals. Monthly autocorrelation might also be important. including this control does not appreciably change the adaptation

Table 6: Robustness to sample and specification changes

	(1) Catch	(2) Catch	(3) Catch	(4) Catch
Niño3.4$_{t-1}$		128.7***	75.1**	71.0**
		(36.9)	(37.7)	(35.0)
Niño3.4$^2_{t-1}$	4.23	-31.8*	9.46	-15.9
	(11.9)	(16.9)	(27.2)	(18.8)
$\widehat{\text{Niño3.4}}_{t-1}$		-104.5***	-112.9***	-79.0**
		(34.4)	(36.7)	(34.4)
$\widehat{\text{Niño3.4}}^2_{t-1}$	-94.6***	-72.9***	-121.1***	-44.6*
	(27.6)	(28.1)	(22.6)	(26.1)
Catch$_{t-1}$				0.41***
				(0.018)
FEs	Baseline	Baseline	Baseline	Baseline
Sample		Latitude< 46°	Drop 1997-2001	
Observations	69,057	46,608	51,920	57,100
R^2	0.078	0.070	0.093	0.22

Notes: The table shows results from estimating versions of equation (12) on monthly data. The dependent variable in each model is the monthly catch, where catch is the number of fish caught. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects unless otherwise noted. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 2 year lags for autocorrelation, unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

effect, although the linear term on the direct effect changes sign.

The revenue calculation is an area where some interpolation was performed to arrive at near-complete observations. This incompleteness comes from two sources First, there is limited geographic coverage in albacore prices. Vessels missing albacore price are simply excluded from the sample when estimating revenue or profit effects since it is unknown by me whether prices in non-U.S. ports follow the same trends as prices in U.S. ports. Among the remaining vessels, 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 12 investigates whether this interpolation procedure is leading to bias in estimates. Column 1 estimates the baseline regression replacing the number of fish caught with the weight of fish

for vessels with daily records of both weight and number of fish. The direct effect of ENSO is slightly higher in this case, but the estimates are, overall, very close to the baseline estimates in Table 2 Column 2. The second column of Table 12 uses interpolated weight as the left-hand-side variable. Results to not change substantially. Finally, column 3 uses revenue with no interpolation, again showing that results are largely unchanged. Overall, these regressions show that the interpolation procedure is not leading to substantive changes in estimates.

As a final robustness check, I want to rule out bias in the forecast coefficient due to variables correlated with expected ENSO but not coming from changes in information. The estimating equation should isolate variation in information by conditioning the forecasts on realizations. The way El Niño and La Niña are announced in the United States offers another way to isolate changes in information. In particular, NOAA declares that an ENSO event is occurring if the Niño 3.4 index is above 0.5 (El Niño) or below -0.5 (La Niña) for 5 consecutive months. This discontinuity in ENSO declaration is unrelated to the physical processes in the ocean, and any realized phenomena caused by ENSO should vary smoothly across the threshold since the Niño 3.4 index is simple a measure of average temperature in the equatorial Pacific, so continuity of Niño 3.4 across the threshold holds.

This result is consistent with harvesters paying particular attention to ENSO around the value at which ENSO events are declared. Any technology or behavior that is always operating, regardless of the Niño 3.4 index value, would not lead to such a jump in output.[19]

6 Adaptation mechanisms

6.1 Adaptation mechanisms conditional on fishing

Table 7 shows estimates for the effect of anticipated and unanticipated changes in ENSO on high frequency decisions of fishing vessels. Each of the outcomes listed in the table are based on daily or intra-trip decisions.

Column 1 of Table 7 shows that if harvesters are able to anticipate a change in ENSO, then they can more accurately target optimal water temperatures, according to the heuristic that fish congregate most in water around 17 or 18°C. The dependent variable in the column is the squared difference between actual water temperature and 17.5°C. In contrast, when the change in ENSO is unanticipated, harvesters are moved

[19]Niño 3.4 exceeding 0.5 is necessary but not sufficient for declaring an El Niño, so in practice, this is a fuzzy regression discontinuity.

Table 7: Intensive margin mechanisms

	(1) Temperature choice error	(2) Hours per day fishing	(3) Fishing lines	(4) Movement extensive	(5) Movement intensive
Niño3.4$_{t-1}$	-0.60	0.099	0.24	6.93	-1.06
	(0.51)	(0.24)	(0.21)	(5.84)	(9.93)
Niño3.4$^2_{t-1}$	0.32	0.22	-0.14	2.84	-9.37
	(0.42)	(0.14)	(0.13)	(3.94)	(6.11)
$\widehat{\text{Niño3.4}}_{t-1}$	1.32*	0.0091	-0.30	-0.51	-14.7**
	(0.70)	(0.27)	(0.24)	(7.35)	(7.50)
$\widehat{\text{Niño3.4}}^2_{t-1}$	-1.74***	-0.64***	-0.31*	-35.4***	4.08
	(0.50)	(0.22)	(0.18)	(8.77)	(4.49)
Average	0.49	11.34	10.39	157.66	1,433.7
Baseline FE	Yes	Yes	Yes	Yes	Yes
Observations	12,430	9,534	12,430	69,057	12,430
R^2	0.095	0.066	0.030	0.062	0.031

Notes: The table shows results from estimating versions of equation (12) 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 2 year lags for autocorrelation. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

further away from the optimal temperature, although this effect is not significant.

Other intensive mechanisms are shown in columns 2 and 3. In response to anticipated extremes in ENSO, harvesters decrease their hours fished per day slightly. The number of lines used per day also appears to go down slightly, although the effect is not strongly significant. The opposite sign responses to realized changes in ENSO for many of these effects point to potential maladaptation *ex post*.

Movement costs and associated net revenue was discussed in Section 5.2. Columns 4 and 5 of Table 7 show that the net revenue improvement estimated in that section is coming largely from changes in extensive margin movement. In other words, vessels are saving fuel costs by sitting out of the albacore fishery.

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. Crew numbers are also fixed. Crew numbers are not observed in the logbook data, and hull length (unsurprisingly) does not change in response to

ENSO. One adaptation that is open to the harvesters on a trip-level frequency and does appear to change with ENSO is the length of the trip, as shown in Table 8.

Table 8: Trip length and frequency

	(1) Fishing days	(2) Transiting days	(3) Trips per month
Niño3.4$_{t-1}$	0.82*	0.063	0.13***
	(0.45)	(0.21)	(0.034)
Niño3.4$^2_{t-1}$	-0.58**	0.014	-0.030*
	(0.24)	(0.12)	(0.015)
$\widehat{\text{Niño3.4}}_{t-1}$	-0.61	0.32	-0.056
	(0.50)	(0.26)	(0.038)
$\widehat{\text{Niño3.4}}^2_{t-1}$	-0.91***	0.13	-0.071***
	(0.32)	(0.15)	(0.022)
Average	11.1	2.42	1.37
Baseline FE	Yes	Yes	Yes
Observations	12,430	4,730	12,430
R^2	0.17	0.024	0.040

Notes: The table shows results from estimating versions of equation (12) 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 2 year lags for autocorrelation. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Column 1 shows that vessels fish slightly fewer days per month given either an expected or unexpected change in ENSO, although the magnitude of the effect is larger in for the expected case. This result is one example of an intensive-margin adaptation that is similar in spirit to entry or exit from the fishery.

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. As indicated by the number of observations, however, transiting days are not recorded for every observation in the dataset.

Finally, trips per month also slightly fall when more extreme ENSO events occur. Like the fishing days result, this decrease in the number of trips comes from both the forecast and realization of ENSO, with the forecast effect being more than twice as large as the realization effect.

6.2 Entry and exit across months

The main results from Table 2 show that much of the adaptation occurring in the fishery is coming from extensive margin adjustments across months. In particular, vessels are choosing to sit out of the albacore fishery during many months of the season rather than risk losses from fish that are too far offshore or that cannot be located. Table 9 looks more closely at this decision.

Table 9: Entry and exit

	(1) Active in the fishery	(2) Exit if active last month
Niño3.4$_{t-1}$	0.049	0.15
	(0.10)	(0.098)
Niño3.4$^2_{t-1}$	0.11*	-0.085
	(0.062)	(0.060)
$\widehat{\text{Niño3.4}}_{t-1}$	0.093	0.17
	(0.12)	(0.15)
$\widehat{\text{Niño3.4}}^2_{t-1}$	-0.53***	-0.088
	(0.10)	(0.11)
Baseline FE	Yes	Yes
Observations	60,695	12,430

Notes: The table shows results from estimating logit model versions of equation (12) on monthly data. The dependent variable in each model is indicated at the top of the column. Additional controls are indicated at the bottom. In parentheses are standard errors clustered at the vessel level. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

The dependent variables in these models are some measures of entry and exit. *Active in fishery* is an indicator equal to one if the vessel is both in the fishery and actively engaged in fishing for albacore. *Exit if active last month* 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 logit models with vessel-level clustering of standard errors.

The entry results show that vessels are much less likely to be active in the fishery if ENSO is forecasted to have extreme values. This result helps explain the drop in output that occurs during extreme ENSO events and also bolsters the movement

results which indicated that most 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. If anything, vessels are slightly more likely to enter the fishery during months with high residual realizations of ENSO, although the effect is small relative to the anticipatory effect.

In contrast, the vessel exit decision is not strongly related to ENSO. This result agrees 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 entry into the fishery a "stickier" state that is not then as responsive to climate shocks.

The vessels are unlikely to be idle during months when they are not actively participating in the albacore fishery. Wise (2011) reports that many fishers also harvest crab and other species during non-albacore-fishing months. Under the assumption that fishing for these other species is not ENSO-sensitive, then welfare calculations based on the adaptation rates calculated in this paper are unaffected.

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 (2013) 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

$$\begin{aligned} y_{it} = \beta_0 + \beta_1 z_{t-1} + \beta_2 z_{t-1}^2 + \beta_3 \hat{z}_{t-1} + \beta_4 \hat{z}_{t-1}^2 + \\ \beta_5 \hat{z}_{t-1} \text{skill}_{t-1} + \beta_6 \hat{z}_{t-1}^2 \text{skill}_{t-1} + \mathbf{x}_{it}' \alpha + \varepsilon_{it} \end{aligned} \tag{15}$$

where the new variable *skill* is a measure of the *ex ante* quality of the forecast. For the 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 skill measures how much uncertainty the forecast resolves. Therefore, if the agent is risk averse, the skill 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 skill is interacted with the forecast terms.

Table 10: Assessing risk aversion

	(1) Catch	(2) Catch
Niño3.4_{t-1}	113.4***	121.3***
	(36.9)	(37.6)
Niño3.4^2_{t-1}	-2.00	-17.8
	(17.7)	(26.3)
$\widehat{\text{Niño3.4}}_{t-1}$	235.1**	-67.7
	(100.3)	(41.7)
$\widehat{\text{Niño3.4}}^2_{t-1}$	-48.0	-150.0***
	(58.2)	(32.2)
$\widehat{\text{Niño3.4}}_{t-1} \times$ skill	-627.2***	
	(180.6)	
$\widehat{\text{Niño3.4}}^2_{t-1} \times$ skill	-222.8**	
	(113.6)	
$\widehat{\text{Niño3.4}}_{t-1} \times$ ens. error		-120.7***
		(32.0)
$\widehat{\text{Niño3.4}}^2_{t-1} \times$ ens. error		108.9***
		(23.3)
Vessel FE	Yes	Yes
Year FE	Yes	Yes
Month FE	Yes	Yes
Observations	69,057	67,715
R^2	0.080	0.081

Notes: The table shows results from estimating equation (15) on monthly data. The dependent variable in each model is total catch per month. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 2 year lags for autocorrelation. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

I measure *ex ante* forecast quality in two ways. First, I calculate the normalized root-mean squared error of the ensemble forecast during the previous two years and normalize that by dividing by the root-mean squared error of a persistence forecast of the Niño 3.4 index. I subtract this normalized value from 1 to create what weather forecasters call the Brier Skill Score (Hamill and Juras, 2006). A value of this measure at 1 means that the forecast is perfectly accurate relative to the naïve persistence forecast. Small or negative numbers mean that the forecast is inaccurate. The skill

measure used in Table 10 is the two-year moving average of this measure for all periods prior to the estimation month, t. One should expect that a risk-averse agent will adapt more if this skill measure is higher.

The second measure of skill is the standard deviation of the forecast plume each month, labeled *ens. error* in the table. Because multiple forecasts are issued beginning in the 1990s, the standard deviation of the plume gives a summary measure of disagreement in the different forecasts. This measure is model-dependent and influenced by model errors, so it does not necessarily represent the full probability distribution of a single forecast (preventing its use as a perfect measure of the second moment of the 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.

Table 10 shows results from estimating Equation (15). The estimates indicate that risk preferences are a potentially important factor in this context. If the skill of the forecast has been higher in recent periods, then agents adapt much more strongly, as shown by the relatively large magnitude of the coefficient on the forecast squared interacted with skill. Similarly, column 2 shows that if the forecast plume is wider, adaptation falls. Both of these results are consistent with preferences for more certain forecasts. The results also show that agents are responding to forecast-specific characteristics, lending support to the assumption that agents are directly consuming these predictions.

7.2 Learning about ENSO and forecasts

By using a single public forecast to measure adaptation, the results assume that all individuals have the same beliefs about ENSO. Differences in ability to understand forecasts, heterogeneity in risk tolerance, or access to private information could alter the conclusions.[20] Here, I focus on heterogeneity in experience with ENSO. A captain or vessel owner with more experience fishing during ENSO conditions might be better equipped to handle the adverse climate, increasing adaptation. In contrast, a captain who has repeatedly suffered from forecasts that missed the realization by a wide margin might be less likely to trust the forecast in the future.

Table 11 investigates this hypothesis in the context of intensive margin catch. Previous results showed that harvesters, on average, had a harder time adapting to ENSO once they had entered the fishery. By including vessel-specific trends that

[20] See, for instance, Kala (2015) for recent evidence on behavioral responses to weather risk.

Table 11: Learning about ENSO

	Catch
Niño3.4$_{t-1}$	-5.00
	(125.0)
Niño3.4$^2_{t-1}$	-129.1**
	(60.2)
Niño3.4$^2_{t-1}$ × ENSO experience	8.25
	(8.10)
$\widehat{\text{Niño3.4}}_{t-1}$	-169.1
	(132.6)
$\widehat{\text{Niño3.4}^2}_{t-1}$	50.6
	(82.2)
$\widehat{\text{Niño3.4}}^2_{t-1}$ × ENSO experience	-56.9**
	(24.2)
V_n(A), 1 event	0.05
	(0.55)
V_n(A), 3 events	0.53**
	(0.24)
V_n(A), 6 events	0.79***
	(0.17)
Baseline FE	Yes
Vessel trend	Yes
Observations	12,430
R^2	0.14

Notes: The table shows results from estimating a modified version of equation (12) on monthly data. The dependent variable is the log of catch, where catch is the average number of fish caught per day in the month. 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 2 year lags for autocorrelation. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1. Standard errors for adaptation are calculated using the delta method.

increment each time a vessel experiences an El Niño or La Niña event, the ability for harvesters to learn can be assessed. Overall, the results suggest that there is an important learning effect. Vessels that have been through more ENSO events adapt at a higher rate. For a novice vessel, adaptation is minimal or non-existent. The average vessel in the dataset has experienced 3 ENSO events, and for this vessel, intensive margin adaptation is moderate. For very experienced vessels—only about 20% of vessels have experience with 6 or more events—intensive margin adaptation is

nearly as effective as total extensive and intensive margin adaptation. Importantly, this adaptation improvement comes both from loading more of the ENSO effect onto the forecast and by reducing the direct effect from ENSO. Results are similar if El Niño and La Niña event are considered separately.

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. Individual and firm adaptation will occur to counter some of these impacts, and if public policy is not appropriately aggressive, such adaptation could bear the brunt of impact reduction. Adaptation does not occur in a vacuum, however. Individuals need to know about their own risks to make informed choices over potential adaptive responses. The importance of this issue makes it crucial to assess the role of information in affecting forward-looking adaptation and 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 functions, this paper assesses the degree of forward-looking adaptation using an estimating equation informed by a 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 ENSO.

From a methodological standpoint, the empirical strategy presented here is not unique to the setting. 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.

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 any bearing on the effects of long-run climate change.

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

On the other hand, if adaptation mechanisms are exhausted, if agents hit corner solutions, if the prices of adaptation mechanisms rise too rapidly, or if climate change causes more extreme weather impacts, then short-run adaptation estimates will not be as good of a guide for the long run. In the setting of this paper, the primary 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 never profitable in expectation, then this adaptation 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.

These results are encouraging for the prospects of adaptation by other highly mobile firms with ready access to non-climate exposed production processes. Caution should be exercised, however, in over-interpreting the results as indicating that these settings will be robust to long-term climate change. Indeed, 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 precarious long-run position—a return to the Malthusian edge.

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 (Grossman and Stiglitz, 1980).

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

Auffhammer, M., S. M. Hsiang, W. Schlenker, and A. Sobel (2013). Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*, 181–198.

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 200211: 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 US Temperature-Mortality Relationship over the 20th Century. *Journal of Political Economy 124*(1), 46.

Brunner, A. D. (2002). El Niño and World Primary Commodity Prices: Warm Water or Hot Air? *Review of Economics and Statistics 84*(1), 176–183.

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

Burke, M., S. M. Hsiang, and E. Miguel (2015). Global non-linear effect of temperature on economic production. *Nature 527*, 235–239.

Butler, E. E. and P. Huybers (2013). 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.

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), 19511958. *Fish Bulletin 115*, 128.

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

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 (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics 3*(4), 152–185.

Deschênes, O. and E. Moretti (2009). Extreme weather events, mortality, and migration. *The Review of Economics and Statistics 91*(4), 659–681.

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.

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

Graff Zivin, J. and M. Neidell (2009). Days of haze: Environmental information disclosure and intertemporal avoidance behavior. *Journal of Environmental Economics and Management 58*(2), 119–128.

Graff Zivin, J. and M. Neidell (2013). Environment, Health, and Human Capital. *Journal of Economic Literature 51*(3), 689–730.

Graff Zivin, J. and M. Neidell (2014). Temperature and the allocation of time: implications for climate change. *Journal of Labor Economics 32*(1), 1–26.

Graff Zivin, J., M. Neidell, and W. Schlenker (2011). Water Quality Violations and Avoidance Behavior: Evidence from Bottled Water Consumption. *American Economic Review Papers & Proceedings 101*(3), 448–453.

Granger, C. W. J. (1969). Prediction with a Generalized Cost of Error Function. *OR 20*(2), 199–207.

Greenstone, M. and J. Gallagher (2008). Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program. *The Quarterly Journal of Economics 123*(3), 951–1003.

Grossman, S. and J. Stiglitz (1980). On the impossibility of informationally efficient markets. *The American Economic Review 70*(3), 393–408.

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.

Hall, R. E. (1978). Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy 86*(6), 971–987.

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. (2016). Climate Econometrics. *NBER Working Paper*, 33.

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.

Inoue, M. and J. O'Brien (1984). A forecasting model for the onset of a major El Niño. *Monthly weather review 112*(11), 2326–2337.

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.

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, nov). 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 oceanbiogeochemical populations dynamics model. *Fisheries Oceanography*, 483–494.

Lemoine, D. and C. Traeger (2014). Watch Your Step: Optimal Policy in a Tipping Climate. *American Economic Journal: Economic Policy 6*(1), 137–166.

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. *The 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.

Neidell, M. (2009). Information, Avoidance Behavior, and Health: The Effect of Ozone on Asthma Hospitalizations. *The Journal of Human Resources 44*(2), 450–478.

Patton, A. J. and A. Timmermann (2012). Forecast Rationality Tests Based on Multi-Horizon Bounds. *Journal of Business & Economic Statistics 30*(1), 1–17.

Reynolds, R. W., N. A. Rayner, T. M. Smith, D. C. Stokes, and W. Wang (2002). An improved in situ and satellite sst analysis for climate. *Journal of Climate 15*(13), 1609–1625.

Rosenzweig, M. and C. R. Udry (2013). Forecasting Profitability. *NBER Working Paper 19334*, 47.

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

Severen, C., C. Costello, and O. Deschênes (2016). A Forward Looking Ricardian Approach: Do Land Markets Capitalize Climate Change Forecasts? *NBER Working Paper 22413*, 46.

Solow, A., R. Adams, K. Bryant, D. Legler, J. O'Brien, B. McCarl, W. Nayda, and R. Weiher (1998). The Value of Improved ENSO Prediction to U.S. Agriculture. *Climatic Change 39*, 47–60.

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. (2015). Adaptation to Climate Change: Historical Evidence from the Indian Monsoon. *Working Paper*, 1–39.

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.

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.

For simplicity, consider a single input model but without the separability assumption. Formally, let the firm solve

$$\max_x \mathbb{E}_{t-1}[\pi_{it}] = p_{1t}\mathbb{E}_{t-1}[f(x_{it}, Z_{it})] - p_{2t}x_{it}. \tag{16}$$

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

$$p_1\mathbb{E}[f_1(x_{it}, Z_{it})] - p_2 = 0 \tag{17}$$

By the implicit function theorem, one can find

$$\frac{\partial x}{\partial \mathbb{E}[Z]} = -\frac{\partial \mathbb{E}[f_1(x,Z)]}{\partial \mathbb{E}[Z]}\left(\frac{\partial \mathbb{E}[f_1(x,Z)]}{\partial x}\right)^{-1} = -\frac{\frac{\partial \mathbb{E}[f_1(x,Z)]}{\partial \mathbb{E}[Z]}}{\mathbb{E}[f_{11}(x,Z)]}. \tag{18}$$

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 nonparametic instrumental variables models with non-separable error.

Let the optimal input choice be

$$x_t^* = \operatorname*{argmax}_x \{p_{1t}\mathbb{E}[f(x_t, Z_t)|\hat{\mathbf{z}}_{t|t-1}] - p_{2t}x_t\}, \tag{19}$$

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 $\mathscr{G}_{t-1}$.[21] This problem yields an optimal

[21] Under loss functions discussed in Section A.5, this vector is simply the conditional expectation of Z_t.

choice for x denoted $x_t^* = h(\hat{\mathbf{z}}_{t|t-1}, \eta_t)$ where η contains everything that shifts factor 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.[22] 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 easily extends to the case of discrete adaptations. Continuity and differentiability does 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

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

The value and normalized value of adaptation can be defined analogously.

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

[22] More details on this can be found below.

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. Given particular functional forms for $\mathbb{E}[g(Z)]$, more instruments could potentially be generated, but there is no guarantee that the number of instruments will equal the number of inputs. More importantly, since 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 since 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

$$\max_{x_{2t}} \pi_t = f(x_{1t}^*, x_{2t})g(z_t) - p_1 x_{1t}^* - p_2 x_{2t} \tag{20}$$

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.

This set-up is easily amenable to dynamic modeling where x_1 is capital and x_2 is consumption or labor. For instance, consider the Euler equation from a standard dynamic, stochastic growth model where C is consumption, X is investment, A is technology, K is capital, and u is the utility function of a representative consumer.

$$E_t\left[\frac{\beta u'(C_{t+1})(1 + A_t f_1'(X_{t+1}, Z_{t+1}))}{u'(C_t)}\right] = 0$$

The particular functional form through which beliefs about the future environmental process enter utility or output will depend on the context and can still result in all adaptation being *ex ante*. For instance, in the Hall (1978) quadratic utility

formulation, consumption in period t is

$$C_t = \left(\frac{r}{1+r}\right)\left(\mathbb{E}_t \sum_{j=0}^{\infty} \left(\frac{1}{1+r}\right)^j Af(Z_{t+j}) + K_t\right)$$

Therefore, consumption is a function of the expected value of the weather process each period in the future.

Empirically decomposing amelioration behaviors and direct effects is challenging in general. Formally, one can think of realizations as unbiased, zero variance forecasts, which allows one to still define all adaptation as "forward looking" in a trivial sense. But, this will lead to a fundamental identification problem since such a forecast cannot be distinguished from weather realizations. Thus, all adaptation estimates based on accurate expectation proxies are, at best, lower bounds on total adaptation in any setting with *ex post* adaptation mechanisms and accurate beliefs about realizations.[23] This issue should not be confused, however, with agents taking actions because weather realizations caused them to update their belief about future weather. In this case, realizations are driving *ex ante* behavior through changes in expectations.

A.4 Adaptation to a non-stochastic environment

The empirical method presented in this paper uses forecasts to identify forward-looking adaptation, so there must be some uncertainty about weather at the time of some of the firm's input decisions for the method to work. One can still gain intuition for the various terms defined in Section 2, however, by examining a non-stochastic version of the firm's decision problem.

Consider a profit maximization problem where the firm chooses an input, x, which enters a production (or revenue) function, $f(x, z)$, which is also a function of weather, z, known at the time of the input decision. Let costs be either linear or convex in inputs and denote them by $c(x)$. Assume that f is at least twice continuously

[23]The need for accurate beliefs about realizations leaves open some possibilities. First, in some forecasting settings, zero-horizon forecasts are issued and do sometimes have errors with respect to realizations that could be exploited. Second, knowing how people learn about something like the weather might shed light on discrepancies between even near-term expectations and realizations. One can think of a poorly calibrated thermometer that is the basis for a firm's use of air conditioning. This thermometer allows the AC to run coincident with the realization of the weather state, but the true weather differs from the inputs to the firm's decision. In this case, however, a researcher would need access to an unbiased thermometer, and one might wonder why the firm did not use the better thermometer.

differentiable in x, at least once continuously differentiable in z, and that costs are at least once continuously differentiable in x. Therefore, the firm's problem is

$$\max_x f(x,z) - c(x)$$

The first order condition is the usual equality between marginal product and marginal cost, $f_1(x,z) = c_1(x)$, and applying the implicit function to this condition gives

$$\frac{\partial x^*}{\partial z} = -\frac{f_{12}(x^*,z)}{f_{11}(x^*,z) - c_{11}(x^*)}. \tag{21}$$

This term is adaptation. Denoting output as y and suppressing arguments of functions from here on, we can write the benefit of adaptation (defined in Equation (4)) as

$$V(A) = \frac{\partial y}{\partial x^*}\frac{\partial x^*}{\partial z} = -f_1\frac{f_{12}}{f_{11} - c_{11}} \tag{22}$$

and divide this by the total derivative of output with respect to weather, $f_1(\partial x^*/\partial z) + f_2$, to get

$$V_n(A) = \frac{-f_1\frac{f_{12}}{f_{11}-c_{11}}}{-f_1\frac{f_{12}}{f_{11}-c_{11}} + f_2} = \frac{f_1 f_{12}}{f_1 f_{12} - f_2(f_{11} - c_{11})} \tag{23}$$

.

This term approaches 1 as the marginal productivity of the input becomes large and is zero if the marginal productivity is zero. The complementarity between inputs and weather acts the same way.

In the case where weather and inputs are multiplicatively separable (the one considered in the body of the paper), signing the normalized benefit of adaptation is simplified. In that case, $\partial^2 y/\partial x \partial z = (\partial y/\partial x)(\partial y/\partial z)$, so the sign of the cross partial derivative will equal the sign of the change in output with respect to weather. The second order condition requires that $f_{11} - c_{11} < 0$. The f_2 term can be canceled out, so the denominator will always be strictly greater than the numerator, and the whole $V_n(A)$ term will be greater than or equal to zero.

A.5 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

$$E[L^p(Z_t, \hat{Z}_t, \mathbf{X}(\hat{\mathbf{Z}})_t, \mathbf{p}_t)|\mathscr{G}_{t-h}] \tag{24}$$

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 (24) 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 $Z_t = 0$ and either of the two following conditions hold

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

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. □

Now, we are ready for the first set of results, which are versions of the forecast sufficiency assumption stated in Section 4. Assume that $\mathscr{G}_t \subseteq \mathscr{F}_t$, or 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 $s^p_{t+h|t} = s_{t+h|t}$. □

We will also be interested in what happens as the public forecast becomes arbitrarily accurate. Define the skill of the forecast as

Definition A.1. The *Brier skill score* or *skill* of a forecast is

$$1 - \frac{MSE}{MSE_C}$$

where MSE is the MSE of the forecast and MSE_C is the MSE of a climatological or reference forecast.

Then a perfectly skillful or accurate forecast has a score of 1.

Now we can show the simple result that if public forecasts are perfectly skillful, then they will provide a perfect proxy for private beliefs.

Corollary A.4. *If the public forecast has perfect* ex ante *skill, then the private expectations equal the public forecast.*

Proof. An MSE-forecast, $s_{t+h|t}$, is unbiased for all h by forecast property 1. Therefore, a forecast will have perfect skill iff $\mathbb{V}(\varepsilon_{t+h|t}) = 0$. Now, assume that $\mathscr{G}_t \supset \mathscr{F}_t$. Then

$$0 \leq \mathbb{V}(\varepsilon^p_{t+h|t}) < \mathbb{V}(\varepsilon_{t+h|t}) = 0,$$

a contradiction. Therefore, $\mathscr{G}_t \subseteq \mathscr{F}_t$, and Prop. A.3 gives the result. □

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, since forecasts based on $\mathscr{F}_t$ are public, the firm will incorporate the public forecast with their private information, leading to $\tilde{s}^p_{t|\tau} = g(s^p_{t|\tau}, s_{t|\tau})$. For instance, with arithmetic mean pooling

$$\begin{aligned} \tilde{s}^p_{t|\tau} &= (1/2)(s^p_{t|\tau} + s_{t|\tau}) \\ \Rightarrow &\frac{\partial \tilde{s}^p}{\partial s} = \frac{1}{2} \end{aligned}$$

which will generally outperform a non-pooled estimator.

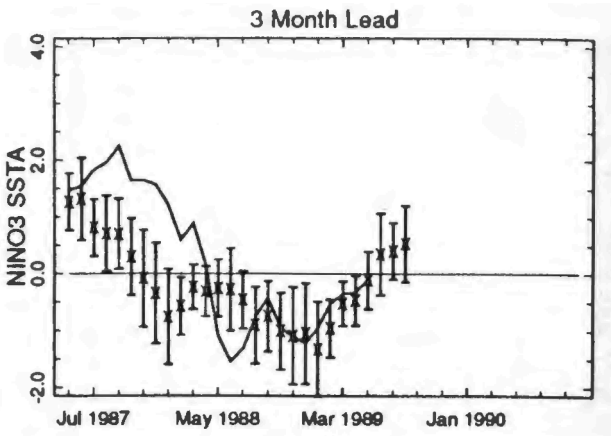
Optimal ensembling by the firm will yield $\mathscr{G}_t \supseteq \mathscr{F}_t$ in all cases where $s_{t+h|t}$ is sufficient for $\mathscr{F}_t$. Therefore, in the event that the public forecasts are not sufficient for the private beliefs of the agent, the ideal estimation strategy would be to instrument for agent beliefs using the public forecasts.

B Data construction details

B.1 ENSO forecast data

Gathering actual contemporary forecast values (what I call "real time" or "historical" forecasts) was central to the project, because accurate knowledge of the information sets available to harvesters is crucial for identification. Unfortunately, to my knowledge, there does not exist a database of real time ENSO forecasts from their initiation in 1989 to the present. Thus, I gathered real time forecasts from the Climate Diagnostics Bulletin (CDB) and the IRI Niño 3.4 summary. The CDB started releasing forecasts in June 1989 and began incorporating the IRI summaries in April 2003. By the year 2000, the number of forecasts incorporated into the Bulletin had grown from 1 to 8.

Figure 2: 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 CDC CCA, LDEO1, LDEO2, LDEO3, LIM, and NCEP forecasts. Other forecasts were either issued as maps or contained idiosyncratic issues that prevented digitization.

For data from 2002 through 2010, I used IRI data helpfully supplied to me by Anthony Barnston. These IRI data have formed the basis for analyses of ENSO forecast performance as 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 Climate Prediction Center Canonical Correlation forecast at a 3 month lead, 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 numerous 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 2010 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 2010 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 2010) 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 Teleconnection

To quantify the relationship between ENSO and temperatures outside of the Niño 3.4 region—what climatologists call teleconnection—I use monthly 1981-2010 satellite measures of sea surface temperatures at a $(1/4)^\circ$ spatial resolution from Reynolds et al. (2002). Temperatures from so called "reconstruction analyses" like this are recommended for use in climate studies by Auffhammer et al. (2013). I define teleconnection as the correlation between temperature in a given location and the Niño 3.4 index from the month prior. I calculate separate teleconnection measures for each month of the year for a given location, reflecting the time-varying strength of ENSO within the year. In particular, ENSO events typically manifest in April or May and last through the beginning of the next year, meaning that effect of ENSO will generally be more apparent in the latter half of the year (Hsiang et al., 2011).

Formally, let m be the month, y be the year, x be the location, and L be a lag length in months. Let $\text{nino}_{m,y}$ be the Niño 3.4 index value for month m in year y, $T_{x,m,y}$ be the temperature at location x, month m, and year y. Let $\rho_{x,m}(L) = \text{corr}(\text{nino}_{m,y}, T_{x,m+L,y})$ for all y. I define teleconnection as this correlation when $L = 1$, or $\text{tel}_{x,m} = \rho_{x,m}(1)$. This definition follows the one used in Hsiang et al. (2011).

The teleconnection value is what is shown in Figures 6, 7, and 8.

B.5 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.6 Catch weight

Catch weight was not recorded in the logbook records for 63,435 of the 193,561 daily records for the full sample (1981 - 2010). 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. Trip weight records were used for interpolation in 11,396 of the missing cases. For the remaining cases, a regression of weight on gear type, year, and month was used to estimate weight.

Table 12: Robustness to interpolation of catch weight

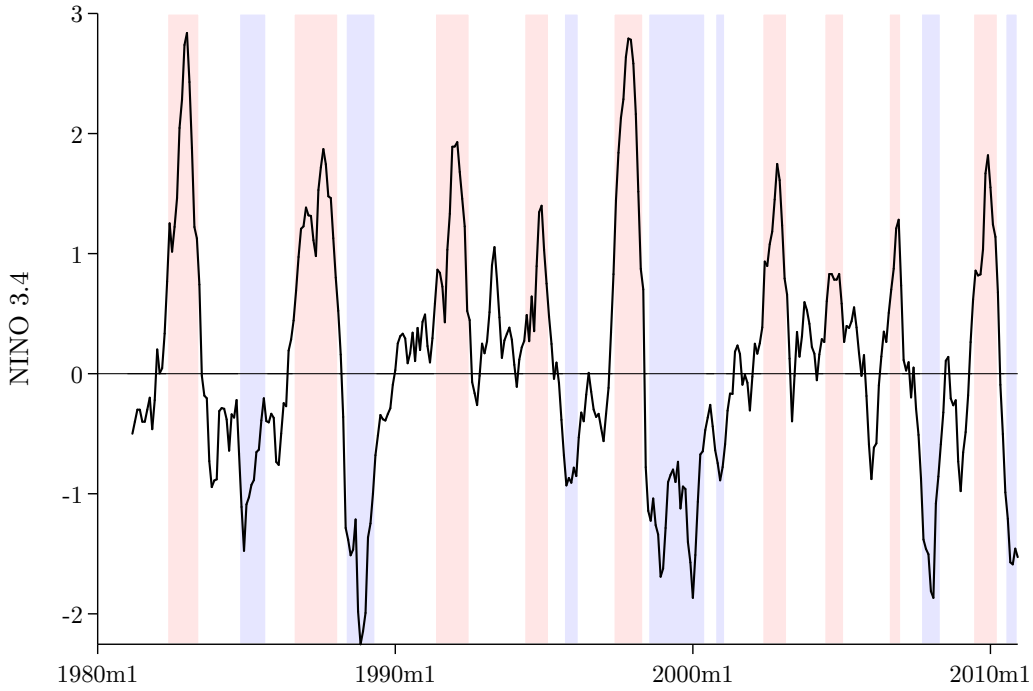
	(1)	(2)	(3)	(4)
	Catch	Catch weight	Catch weight, interpolated	Revenue
Niño3.4_{t-1}	95.9***	1269.1***	1542.6***	1204.6***
	(35.2)	(477.1)	(529.8)	(451.5)
Niño3.4^2_{t-1}	-14.2	-175.3	-222.0	-66.9
	(15.1)	(207.0)	(227.2)	(191.4)
$\widehat{\text{Niño3.4}}_{t-1}$	-91.6***	-1250.6***	-1492.5***	-1433.6***
	(30.2)	(402.0)	(452.9)	(356.3)
$\widehat{\text{Niño3.4}}^2_{t-1}$	-69.3***	-1006.0***	-1075.1***	-771.7**
	(26.3)	(348.2)	(385.0)	(314.5)
FEs	Baseline	Baseline	Baseline	Baseline
Weight measure	Observed	Observed	Interpolated	Observed
Observations	67,232	67,232	69,057	67,232
R^2	0.073	0.066	0.076	0.066

Notes: The table shows results from estimating versions of equation (12) on monthly data. The dependent variable in each model is the monthly catch, where catch is the number of fish caught. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects unless otherwise noted. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 2 year lags for autocorrelation, unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Table 12 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 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. Finally, 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.

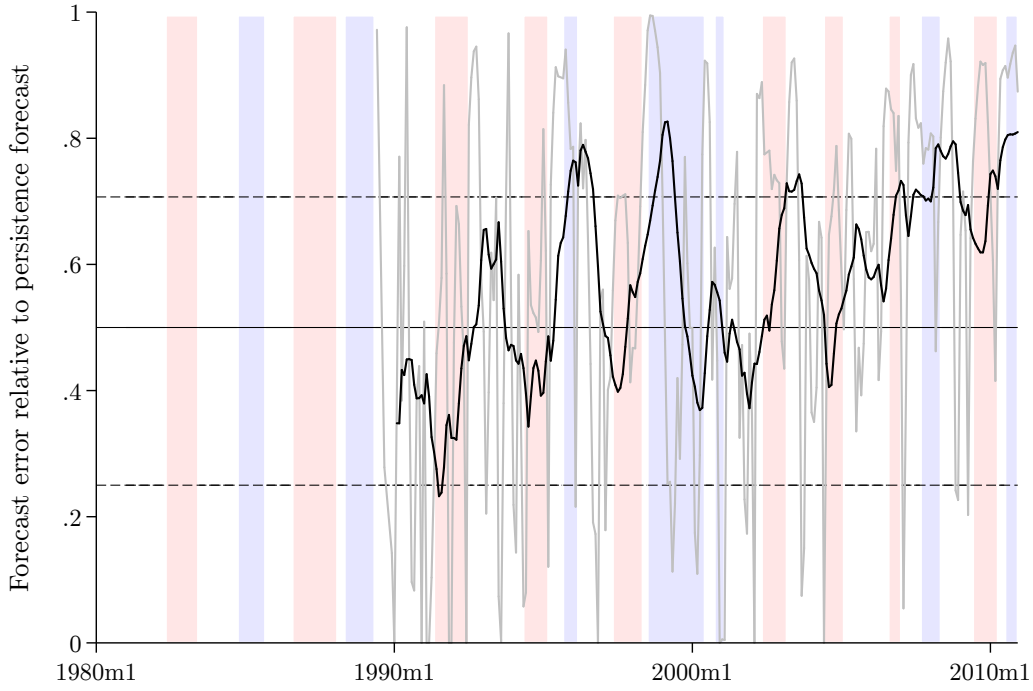
C Additional figures and tables

Figure 3: ENSO Cycle



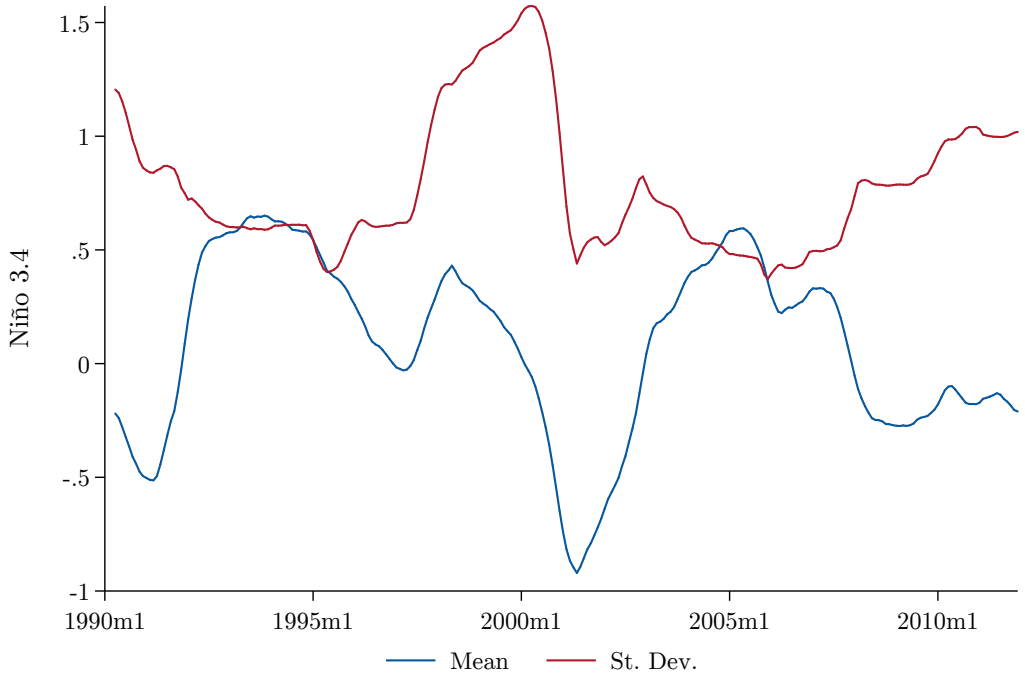
Notes: The ENSO cycle is represented here by the NINO3.4 index, which is the three month moving average of SST anomalies from the NINO3.4 region of the Pacific. Values above 0.5 indicate an El Niño and values below -0.5 indicate La Niña, as denoted by the red and blue shaded regions respectively. For more information on this series, see Section 3.

Figure 4: Forecast skill



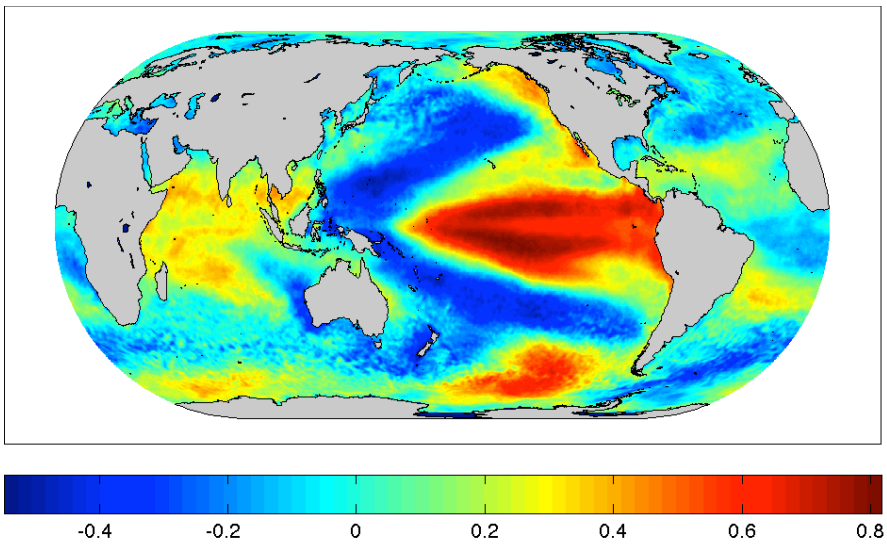
Notes: Forecast skill is indicated by the light gray lines, and the 12 month moving average of skill is given by the blue lines. Skill is the rolling mean squared error of forecasts normalized by the rolling mean squared error of a naïve persistence forecast. For details, see Section A.5. El Niño periods are indicated in red, and La Niña periods are indicated in blue.

Figure 5: Moving standard deviation of ENSO



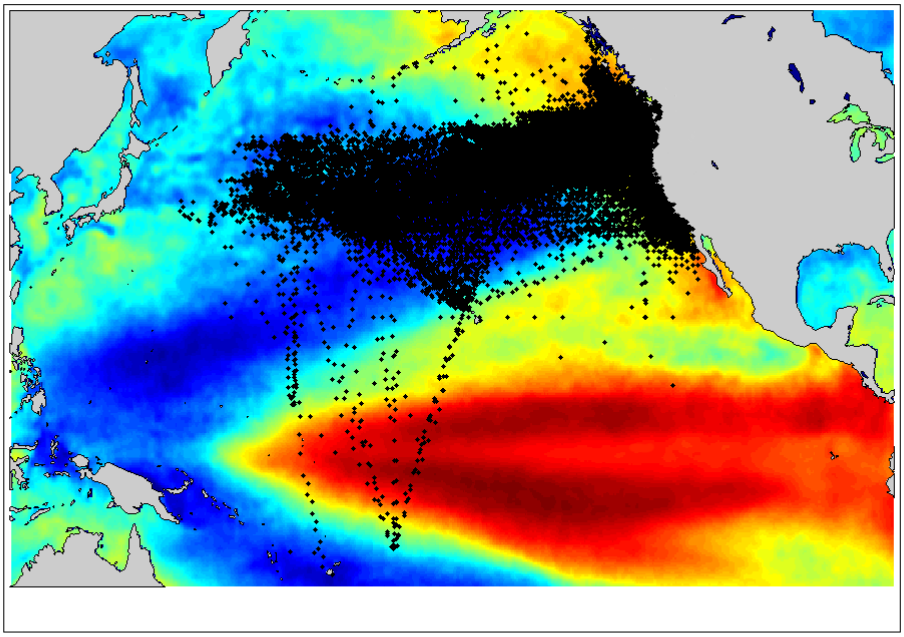
Notes: Moving average and standard deviation of the Niño 3.4 index is shown for the main estimation sample. Rolling values use a three year window and monthly data.

Figure 6: Teleconnection between Niño 3.4 and sea surface temperature



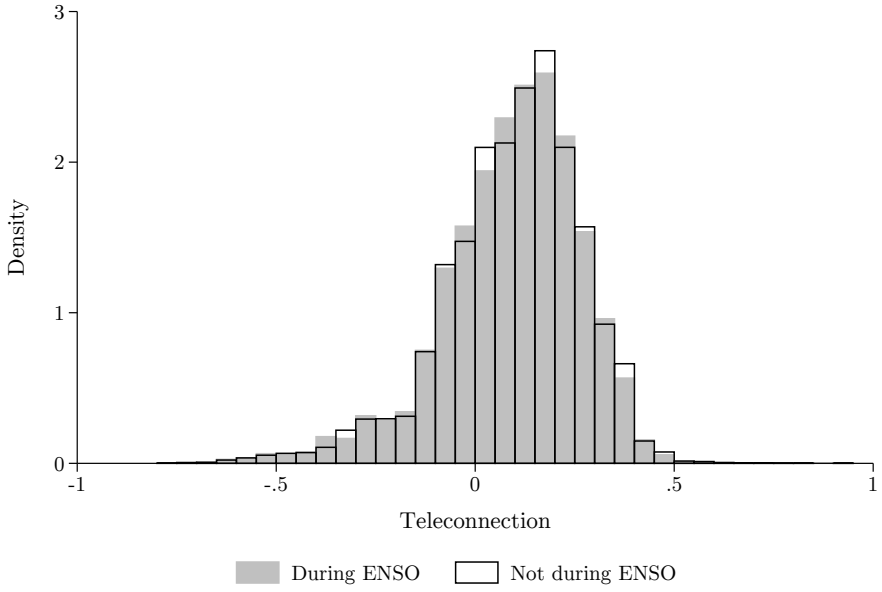
Notes: The heat map shows correlation between the one month lag of the Niño 3.4 index and sea surface temperature for each quarter degree latitude-longitude grid cell. This correlation serves as the teleconnection measure in this paper. For more information on this calculation, see Section 3.

Figure 7: Fishing locations across the North Pacific



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 6. Each point shows a daily observation of either fishing or transiting.

Figure 8: Teleconnection during ENSO events versus not



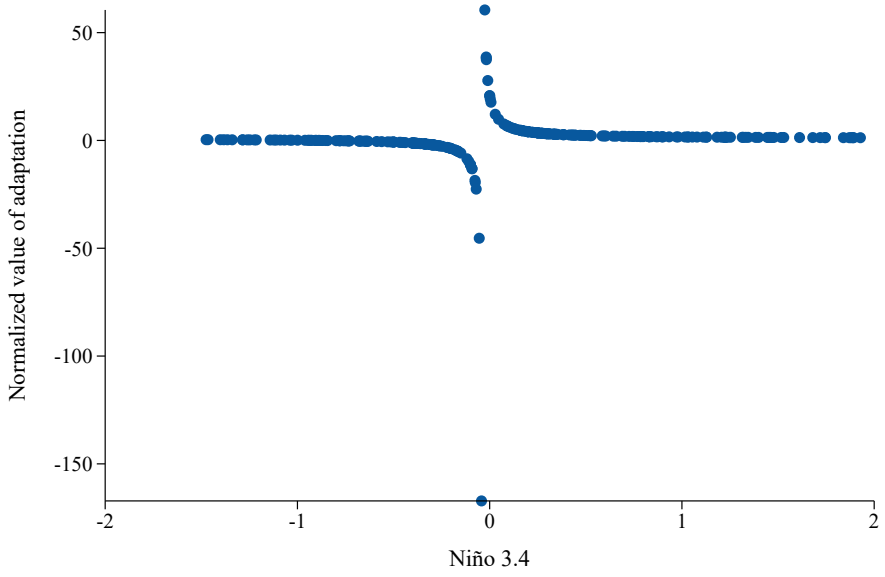
Notes: Two histograms of daily teleconnection status are shown. The gray is during ENSO events, and the black outline is not during ENSO events.

Table 13: Additional robustness

	(1) Catch	(2) Catch
Niño3.4_{t-1}	111.8***	123.8**
	(41.7)	(52.7)
Niño3.4^2_{t-1}	-5.72	-11.9
	(19.1)	(24.4)
$\widehat{\text{Niño3.4}}_{t-1}$	-81.0*	-84.4*
	(45.7)	(44.8)
$\widehat{\text{Niño3.4}}^2_{t-1}$	-72.0**	-69.9**
	(31.3)	(30.0)
$\widehat{\text{Niño3.4}}_{t}$	-30.6	-8.14
	(61.9)	(57.9)
$\widehat{\text{Niño3.4}}^2_{t}$	-14.3	-27.1
	(36.6)	(31.9)
$\widehat{\text{Niño3.4}}_{t+1}$		-34.5
		(72.0)
$\widehat{\text{Niño3.4}}^2_{t+1}$		19.6
		(32.1)
FEs	Baseline	Baseline
Observations	67,715	67,260
R^2	0.079	0.079

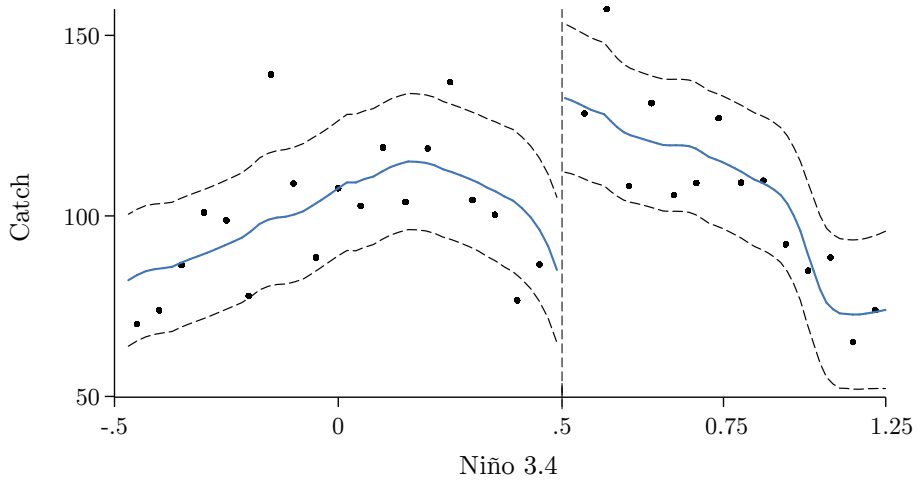
Notes: The table shows results from estimating versions of equation (12) on monthly data. The dependent variable in each model is the monthly catch, where catch is the number of fish caught. In addition to the listed variables, all models contain vessel, year, and month-of-year fixed effects unless otherwise noted. In parentheses are spatial-temporal HAC robust standard errors using a uniform kernel, a distance cutoff of 30km, and 2 year lags for autocorrelation, unless otherwise noted. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Figure 9: Normalized value of adaptation as a function of Niño 3.4 values



Notes: The normalized value of adaptation, $V_n(A)$ is shown for the revenue estimates in Table 2. For details on the calculation of this value, see Section 5.3.

Figure 10: Regression discontinuity of catch with respect to Niño 3.4



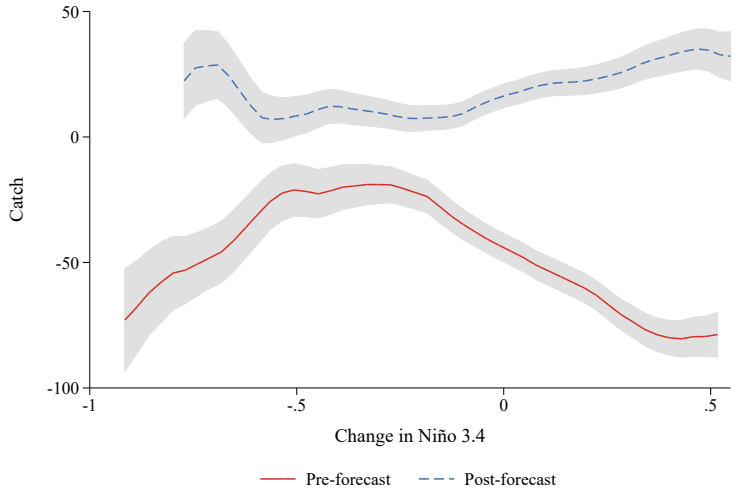
Notes: Each point is the average catch in 0.05° bins of the Niño 3.4. Local linear regressions (Epanechnikov kernel with bandwidth of 0.1) are fit to the data that fall on either side of Niño 3.4= 0.5, the pre-requisite for declaring an El Niño.

Table 14: Price effects of ENSO

	Albacore price	Fuel price
Niño 3.4	-0.069	-0.13*
	(0.070)	(0.071)
Niño 3.4^2	-0.039	-0.014
	(0.053)	(0.045)
Constant	0.97***	1.99***
	(0.074)	(0.076)
Observations	31	347

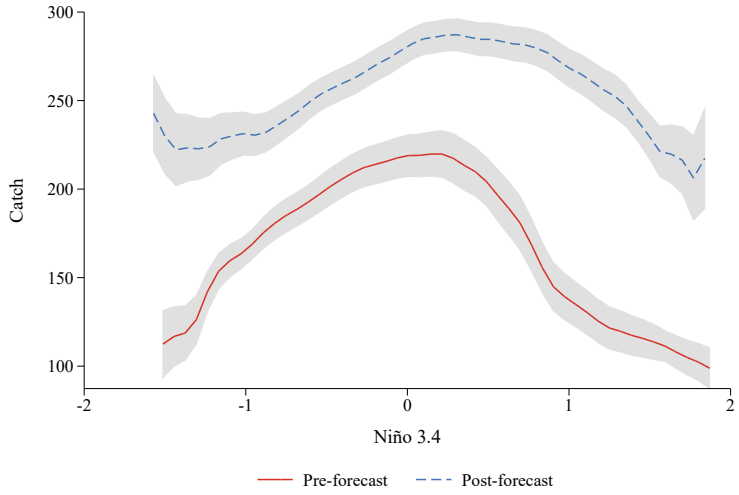
Notes: The table shows results from estimating Newey-West regressions on monthly (fuel prices) or annual (albacore prices) data. The dependent variable in each model is indicated at the top of the column. In parentheses are Newey-West standard errors with 2 lags for autocorrelation. Significance indicated by: *** p<0.01, ** p<0.05, * p<0.1.

Figure 11: Output and ENSO before and after forecasts, changes



Notes: Each line shows a local linear regression (Epanechnikov kernel with bandwidth of 0.18) of catch on the change in the Niño 3.4 index between month $t-1$ and $t-2$. All variables are residualized on month. The red, solid line uses the sample from 1981 to May 1989 before ENSO forecasts were released. The blue, dashed line uses the sample from after forecasts were released in June 1989 until 2010. Shaded areas give the 95% confidence intervals.

Figure 12: Output and ENSO before and after forecasts, raw data



Notes: Each line shows a local linear regression (Epanechnikov kernel with bandwidth of 0.38) of catch on the Niño 3.4 index the previous month. The red, solid line uses the sample from 1981 to May 1989 before ENSO forecasts were released. The blue, dashed line uses the sample from after forecasts were released in June 1989 until 2010. Shaded areas give the 95% confidence intervals.