EPJ manuscript No.
(will be inserted by the editor)

Forecasting in the presence of expectations

Roy Allen[1], Joshua Graff Zivin[1,2], and Jeffrey Shrader[1,a]

[1] Department of Economics, University of California, San Diego, 9500 Gilman Drive, San Diego, CA 92093

[2] School of International Relations and Pacific Studies (IR/PS), University of California, San Diego, 9500 Gilman Drive, San Diego, CA 92093-0519

Abstract. Physical processes routinely influence economic outcomes, and actions by economic agents can, in turn, influence physical processes. This feedback creates challenges for forecasting and inference, creating the potential for complementarity between models from different academic disciplines. Using as example of prediction of water availability during a drought, we illustrate the potential biases in forecasts that only take part of a coupled system into account. In particular, we show that forecasts can alter the feedbacks between supply and demand, leading to inaccurate prediction about future states of the system. Although the example is specific to drought, the problem of feedback between expectations and forecast quality is not isolated to the particular model—it is relevant to areas as diverse as population assessments for conservation, balancing the electrical grid, and setting macroeconomic policy.

1 Introduction

The interactions between human and natural systems play an important role in determining environmental and economic outcomes. Environmental conditions determine the scarcity or abundance of resources, which in turn influence the incentives to undertake a variety of economic activities. Those human behaviors also directly impact environmental conditions creating important feedbacks that can hasten or temper environmental impacts. While considerable scholarship has been devoted to research at the nexus of coupled natural and human systems, surprisingly little work has focused on the role played by forecasts in shaping that relationship.

Forecasts are important for at least two reasons. First, they play an increasingly pivotal role in public policy, shaping decision making in areas ranging from transportation and energy to the health and leisure sectors. Second, they are largely designed to influence human activity. Air quality forecasts are meant to influence how much time is spent outdoors, traffic forecasts are designed to influence choices about commute modes and times, and it is hoped that climate forecasts will help us choose appropriate carbon emissions scenarios.

From a modeling perspective, the noteworthy feature of these forecasts is that they can lead to changes in behavior today based on revised expectations about

[a] e-mail: `jgshrade@ucsd.edu`

the future. A forecaster who does not take this behavioral response into account is assuming that individuals, when given the forecast, will continue to act in the same way as they did before the release of the forecast. This is unlikely to be the case, however, since the very act of publicly releasing a forecast increases the likelihood that it will miss its mark due to the attending changes in behaviors that result from altered expectations. Moreover, the solution is not simply to increase the accuracy of the forecast. Ironically, forecasts that are believed to be more accurate can generate larger changes in expectations leading to more powerful feedback within the system, further worsening the error.

To be clear, these forecast feedbacks are distinct from the sorts of economic feedbacks that are often captured in sophisticated modeling exercises. Those feedbacks are based on expected behaviors in response to the *realization* of environmental states, but forecasts now require us to include changes in behavior based on *expectations* regarding yet unrealized environmental states. Ignoring this latter feedback can generate errors in either direction. If the "naïve" forecast predicts that the state of the system will be low in the future, forward looking agents with well defined property rights may use less of the resource today in order to conserve it for use in the future. The reduced use, however, leaves more of the resource unexploited, leading the forecast to exaggerate near-term scarcity and thus appear overly pessimistic. In contrast, if the forecast predicts that the state of the system will be high, agents may exploit the resource more intensively under the false assumption that ample supplies will remain in the future, leading to an overestimate of resource abundance. The presence of public goods and externalities, where property rights are not well defined, may change this calculus. In that case, a low forecast provides each agent with an incentive to rush to use the resource before it is depleted by others, precipitating a race to the bottom.

The form of feedback we consider is distinct from the typical concerns about which variables to include in a statistical forecasting model. As our model illustrates, even forecasters who have proxy variables for each component of the true model and have had good performance with those proxies can routinely issue ex post-inaccurate forecasts if agent expectations lead to changes in the state. Therefore, the failure is not one of excluding important variables but of not taking appropriate causal channels into account. In the context of our model, the forecast *itself* influences the system.

The fundamental challenge to modeling these forecast feedbacks is our limited understanding of expectation formation by forward-looking economic agents and how those expectations link back to the systems being forecast [19]. In this paper, we develop a simple model that provides one example of the problems that can arise when trying to make accurate forecasts in the presence of feedbacks. We look at a stylized example of farmers who make planting decisions based on expectations about water availability that are influenced by forecasts. We show that if forecasts affect the expectations of the agents, those expectations determine actions that the agent takes in advance of the forecasted event being realized, and if those actions influence the event being forecasted, then the forecast will be inaccurate, in potentially complex ways. For clarity, we present a simple model where each part of the system directly influences the other part, but the result can hold more generally and for systems where the causal chain is longer. This model also provides some insight into how to study this process empirically. When analyzing situations where expectations are important, the major challenge for an empirical researcher is measuring those expectations. If agents use forecasts, at least in part, to form their expectations, then forecasts can be used to identify changes in expectations—a method that has begun to be employed in economics research [15,17].

2 Example: Forecasting water levels

We now turn to a stylized model of water level forecasting to highlight the key conceptual point that even with good understanding of a physical system, feedback can lead to erroneous forecasts. These errors can be eliminated if the feedback is correctly taken into account. The topic is timely, with the western United States experiencing a severe and prolonged drought. Under continued climate change, such droughts are expected to increase in frequency for many areas of the country [5]. This example is also not purely theoretical: the National Oceanic and Atmospheric Administration (NOAA), among other groups, regularly release forecasts of drought, groundwater levels, and other hydrological parameters of interest to farmers and the broader public [16].

The system we model involves farmers who need to make judgments about the future supply of water in order to make planting decisions. These decisions, in turn, affect the supply of water since more intensive planting will consume more water (in an area like California, 80% or more of fresh water is used for agriculture [20]). We consider what happens when a forecaster tries to predict the water level, both with and without taking these expectations into account.

The model shows that when farmers respond to forecasts of the water level when making their planting decisions, they influence the realized amount of water. Thus, if a forecaster says that rainfall or the water level will be low, more farmers will decide to forgo planting this season, reducing demand for water and therefore supporting the water level. This will cause the forecast to appear to be erroneously low. Similarly, a forecaster who says that rainfall or the water level will be high will cause more farmers to plant, raising demand for water and placing strain on the water supply. This will cause the forecast to, ultimately, be erroneously high. A forecaster who incorporates the effect he or she has on demand can solve for the optimal prediction that will result in a forecast that matches the realization.

2.1 Farmers

Consider farmers who are each deciding whether to plant crops in year t. At the beginning of the season, farmer i can plant crops, and if the water level during the season, W_t, is sufficiently high, then the crops will grow and can be sold in the market for price p_A. Each farmer can plant the quantity A crops per season, and for simplicity, we assume that the farmer cannot choose to plant just a fraction of his or her land and that the water level must be sufficient for all planted area or else the entire crop fails.[1] Thus, revenue conditional on successful harvest is $p_A A$. It costs the farmer K dollars to plant the crops (for seed, labor, capital rental, and other inputs), which cannot be recouped if the harvest fails. We assume that farmers have idiosyncratic water needs, and the crops of farmer i grow and can be sold when $W_t > \gamma + \eta_i$. Combining these assumptions, if the farmer plants, the expected profit will be

$$\mathrm{E}[y_{it}] = \mathrm{P}_i(W_t > \gamma + \eta_i)(p_A A - K) - \mathrm{P}_i(W_t \leq \gamma + \eta_i)K. \tag{1}$$

$P_i(W_t > \gamma + \eta_i)$ reflects the beliefs of farmer i about the probability his or her crops will succeed. The parameter γ can be interpreted as the average drought susceptibility of the crops. If the farmer plants nothing, we assume, without loss of generality, that

[1] The assumption that all crops fail is made purely for analytical convenience. The conclusions would be substantially unchanged with fractional failure proportional to the water shortfall. Indeed, across multiple runs of the simulation presented below, average crop failure is equivalent to this more general model.

expected profit is zero. Assume that the farmer only plants if he or she expects profits, as defined in (1), to be positive, so planting occurs when $\mathrm{P}_i(W_t > \gamma + \eta_i) \geq K/(p_A A)$.

In principle, we allow each farmer i to assign a different value to $\mathrm{P}_i(W_t > \gamma + \eta_i)$, which can arise due to differences in information or beliefs as well as idiosyncratic water needs. In the spirit of recent work on persuasion of Bayesian agents [11], all farmers have access to a public forecast $\hat{W}_t$ of the predicted amount of water available for use this season, and we assume that each farmer makes decisions based on this public forecast, a personal belief about the state, knowledge of the distribution of η across farmers, knowledge of the structural equations governing demand, and some knowledge about the structure of the public forecasts. The exact level of structural knowledge will be different for each forecasting scenario and will be specified in Section 2.3.

The farmer combines the public forecast and structural knowledge to form a private forecast, $\hat{W}_t^f$, each period, where the superscript f denotes the farmer's forecast. We specify the fraction of farmers who plant, conditional on the forecast $\hat{W}_t^f$, as,

$$\begin{aligned} f(\hat{W}_t^f) &= \mathrm{P}\left(\mathrm{P}_i(W_t > \gamma + \eta_i) \geq K/(p_A A) \mid \hat{W}_t^f\right) \\ &= \begin{cases} 0 & \text{if } \frac{1}{2b}(\hat{W}_t^f - \gamma) - \frac{K}{p_A A} < -1 \\ 1 + \frac{1}{2b}(\hat{W}_t^f - \gamma) - \frac{K}{p_A A} & \text{if } -1 \leq \frac{1}{2b}(\hat{W}_t^f - \gamma) - \frac{K}{p_A A} \leq 0 \\ 1 & \text{otherwise.} \end{cases} \end{aligned} \tag{2}$$

This fraction can be strictly between 0 and 1 due to heterogeneous water needs as reflected in η_i, as well as idiosyncratic beliefs. The parameter b can be thought of as a parameterization of heterogeneity. As $b \to 0$, the fraction who plant tends to either 0 or 1 depending on whether $\hat{W}_t^f > \gamma$. Importantly for our analysis, the fraction planting is increasing in $\hat{W}_t^f$, which reflects that when the private forecast water level is higher, more farmers plant. The assumed relationship between the public forecast $\hat{W}_t$ and $\hat{W}_t^f$ will be made explicit in Section 2.3 as we consider different forecasting scenarios.

Collapsing all public information into a single signal $\hat{W}_t^f$ abstracts from several issues. Knowing the distribution of η, each farmer knows this demand system and can use it to create improved forecasts relative to the forecasters discussed below. Even still, the farmers in this model are highly simplified. For instance, we assume that there is only one public forecast consumed by each farmer. In part, private beliefs proxy for any idiosyncratic forecasts that farmers might be receiving, but in practice, sophisticated farmers might draw on many forecasts of stochastic processes that influence their profit. We discuss extensions to the model along these lines in Section 2.5. Recent research, however, also indicates that some farmers under-utilize forecasts in making early season investment decisions [15]. A more complete analysis would involve describing how and why the farmers might combine their idiosyncratic information [3] as well as the objectives of the forecasters [6]. For example, the forecasters may be more interested in inducing a response from the agents (such as water conservation) than in providing accurate forecasts.

2.2 Water

We are modeling an environment where all crops are watered by irrigation, so the supply of water from rain is added solely to the stock of water from the previous period. No farms receive direct rainfall. Suppose that annual rainfall, denoted R_t, arrives according to

$$R_t = \zeta + \xi \mathbb{1}\{t \geq \tau\} + \epsilon_t, \tag{3}$$

where $\epsilon_t \sim \mathrm{N}(0, \sigma_\epsilon^2)$.[2] The term $\xi \mathbb{1}\{t \geq \tau\}$ denotes a structural change that occurs at time τ and shifts the average rainfall by $\xi < 0$ for every period thereafter, where $\mathbb{1}$ is the indicator function. One can interpret τ as representing the date of the start of a persistent drought. We consider only mean shifts in the rainfall process, which is without loss of generality in this framework, but more complicated models might attend to higher order moments of the rainfall distribution that could change during drought or other weather extremes. For solution stability, we assume $\zeta < 1$ and $\zeta + \xi \geq 0$.

We further assume that time-varying demand is only due to the farmers and is just equal to the fraction of farmers who plant: $D_t = f(\hat{W}_t^f)$. Thus, the water stock evolves according to supply and demand each period.

$$W_t = W_{t-1} + R_t - D_t. \tag{4}$$

2.3 Forecasts

Because farmers use the forecasts when making their planting decisions, the forecasting process plays a crucial role in determining the water level. We therefore consider a variety of forecasts to examine their effects on the system. At the beginning of each period, forecaster j makes prediction $\hat{W}_{jt}$ of water level at the end of time period t. We will consider four types of forecasts: a "statistical" forecaster ($j = a$) who uses standard time series tools to forecast the water level without any knowledge of the underlying structure, an "economic" forecaster ($j = b$) who knows the demand process but not the rainfall process, a "meteorological" forecaster ($j = c$) who knows the rainfall process but not the demand process, and a fully structural forecaster ($j = d$) who knows both the rainfall and demand processes.

The statistical forecaster has noticed that the evolution of W_t is well-approximated by an autoregressive (AR) model of order 1, meaning that the forecaster estimates the model

$$W_{a,t} = \alpha_{a,t} + \rho_{a,t} W_{a,t-1} + u_{a,t} \tag{5}$$

assuming $u_{a,t}$ is independent across t [8]. Here, α and ρ are coefficients to be estimated. The forecaster updates the coefficients according to

$$(\hat{\alpha}_{a,t}, \hat{\rho}_{a,t}) = \operatorname*{argmin}_{\alpha,\rho \in \mathbb{R}^2} \sum_{n=1}^{t-1} (W_{a,n} - \alpha - \rho W_{a,n-1})^2 K_{h_1}(t-n) \tag{6}$$

where $K_{h_1}(t-n)$ is a kernel that down-weights past observations of data. For simulation, we will use a commonly chosen exponentially decreasing kernel $K_{h_1}(t-n) = e^{(-h_1(t-n))}$. Thus, forecasts are based on the historical record of rainfall, with greater weight placed on recent observations.

We will further assume that in this case, like the forecaster, the farmer does not know the equation governing water supply. This assumption is made purely for illustrative purposes so that the effect of fully versus partially sophisticated agents can be explored. Because they do not know the water supply equation, farmers in this case cannot improve on the forecast given to them by the forecaster, so $\hat{W}_{a,t}^f = \hat{W}_{a,t}$.[3]

[2] Negative R_t can be interpreted as loss due to evaporation or a lack of rainfall relative to exogenous urban demand, for instance. Forcing $R_t > 0$ can be done with some technical complications that obscure our analysis.

[3] Past water supply is publicly observable, so it might be possible to improve on the forecast by using an alternative statistical model, but our point is not to compare alternative statistical forecasting methods, so we will ignore that point here.

If farmers had knowledge of the water supply equation, they could filter historical rainfall from the water level and make their own inference on the future path of rain, resulting in private forecasts equivalent to those issued by the economic forecaster, to whom we turn next.

The economic forecaster has an exact model for the demand of the farmers and knows the supply equation 4. This forecaster does not, however, know the physical process governing rainfall and therefore must infer future rainfall from publicly observable, historical rainfall. Knowing supply and demand, the economic forecaster models the evolution of water as

$$W_{b,t} = \alpha_{b,t} + W_{b,t-1} - f(\hat{W}_{b,t}) + u_{b,t} \tag{7}$$

where $\alpha_{b,t}$ is an unknown parameter to be estimated and $f(\hat{W}_{b,t})$ is given by (2). The economic forecaster takes into account how the forecast itself will impact demand. Specifically, given an estimate $\hat{\alpha}_{b,t}$ the economic forecaster chooses $\hat{W}_{b,t}$ to solve

$$\hat{W}_{b,t} + f(\hat{W}_{b,t}) = W_{b,t-1} + \hat{\alpha}_{b,t}. \tag{8}$$

The forecaster updates $\hat{\alpha}_{b,t}$ according to

$$\hat{\alpha}_{b,t} = \operatorname*{argmin}_{\alpha\in\mathbb{R}} \sum_{n=1}^{t-1} (W_{b,n} - \alpha - W_{b,n-1} + f(\hat{W}_{b,n}))^2 K_{h_2}(t-n). \tag{9}$$

We will use the same exponential kernel as in the statistical case, but with parameter h_2 chosen to match first period root mean squared error (RMSE) between the AR(1) statistical forecast and the economic forecast for convenience of interpretation. In this case, the forecaster possess the same information as the farmers, so again the private forecast will coincide exactly with the public forecast, assuming that the same statistical procedure is used to estimate rainfall.

The meteorological forecaster has a structural model for the physical side of this system, so they know, up to noise, what will happen with the rainfall process this season. Formally, they can predict the non-stochastic portion of rainfall, $\hat{R}_t = \zeta + \xi 1\{t \geq \tau\}$, with certainty each period. This forecaster knows, however, that demand influences the ultimate water level, so he or she tries to infer the missing component by estimating

$$W_{c,t} = \alpha_{c,t} + W_{c,t-1} + \hat{R}_t + u_{c,t} \tag{10}$$

by performing the following optimization each period

$$\hat{\alpha}_{c,t} = \operatorname*{argmin}_{\alpha\in\mathbb{R}} \sum_{n=1}^{t-1} (W_{c,n} - \alpha - W_{c,n-1} - \hat{R}_n)^2 K_{h_2}(t-n). \tag{11}$$

For a real-world example of this kind of correction to an otherwise structural (or, in the terminology of climatologists, dynamical) model, see [4]. In this case, the public and private forecasts will not coincide. Assuming that the meteorologist releases both a water level forecast and a rainfall forecast, the farmers can use the rainfall forecast and knowledge of the demand system to create optimal forecasts in a manner equivalent to the structural forecaster discussed next.

Finally, by taking both the expectations of the farmers and the rainfall process into account, a fully structural forecaster can produce a forecast that is exact up to the irreducible noise in the rainfall process (ϵ_t). This forecaster knows that the true process is

$$W_t = W_{t-1} + \hat{R}_t + \epsilon_t - f(\hat{W}_t)$$

and chooses $\hat{W}_{d,t}$ that solves,

$$\hat{W}_{d,t} + f(\hat{W}_{d,t}) = W_{d,t-1} + \hat{R}_t,$$

yielding the same expression as in the case of the economic forecaster, but without the need for inference on the rainfall process. This forecast will be used as a baseline when comparing the quality of the three forecasters discussed above and provides the forecast, $\hat{W}^f_{c,t}$, used by farmers in the case with meteorology forecasts.

2.4 Forecasting during a drought

We now consider how each forecaster performs when the physical system experiences a drought. To perform these analyses, we simulate the model using the parameter values given in Table 1.

Table 1. Simulation parameter values

parameter	value
p_A	2
A	1
K	1
γ	0.4
b	0.5
ζ	0.5
h_1	1/3
h_2	1/2

We examine a permanent drought that occurs after ten periods. In terms of Equation (3), at time $\tau = 10$ this drought reduces rainfall by 30% ($\xi = -0.15$), on average, each period. Figure 1 shows the forecasts of water level (dashed line) and actual water level (solid line) each period for the four types of forecasters discussed above. In each panel, the optimal forecast is shown with asterisks so that the forecasts can be compared to the ideal.

Prior to the drought, the system is in equilibrium with the water level stable at 0.4, the profitability break-even point for farmers. The drop in rainfall should induce some farmers to exit the market once they realize that they will not be able to provide water for their crops. Indeed, one can see that in the fully structural case, Figure 1 panel (d), the system reaches a new equilibrium with a lower overall water level after about 5 periods. This lower water level also reflects lower total farming. With the optimal forecast, the forecaster is able to accurately predict the decline in water use each period and, in fact, smooth the transition between drought and non-drought states—a service that is potentially valuable to firms. This smoothness occurs because water usage by farmers is appropriately conditioned by forecasts in anticipation of future rainfall events.

The other cases each illustrate one or more forecasting failures. The statistical forecaster, panel (a), shows the impact of a structural break (the change in the rainfall process) coupled with incomplete knowledge of the underlying system. The forecaster has some inertia in the forecast, so even after the rainfall shock occurs, this forecast continues to predict a high water level, near the historical average. This optimistic forecast causes too many farms to stay in business so that total water demand far

The figure shows simulated forecasts and water levels in response to a rainfall shock for the four types of forecasters discussed in Section 2.3, with (a) showing the statistical forecast, (b) showing the forecast with known demand, (c) showing the forecast with known rainfall, and (d) showing the optimal forecast where both demand and rainfall are known. The dashed black line in each panel shows the forecast, the solid line shows the water level, and the asterisks show the optimal forecast.

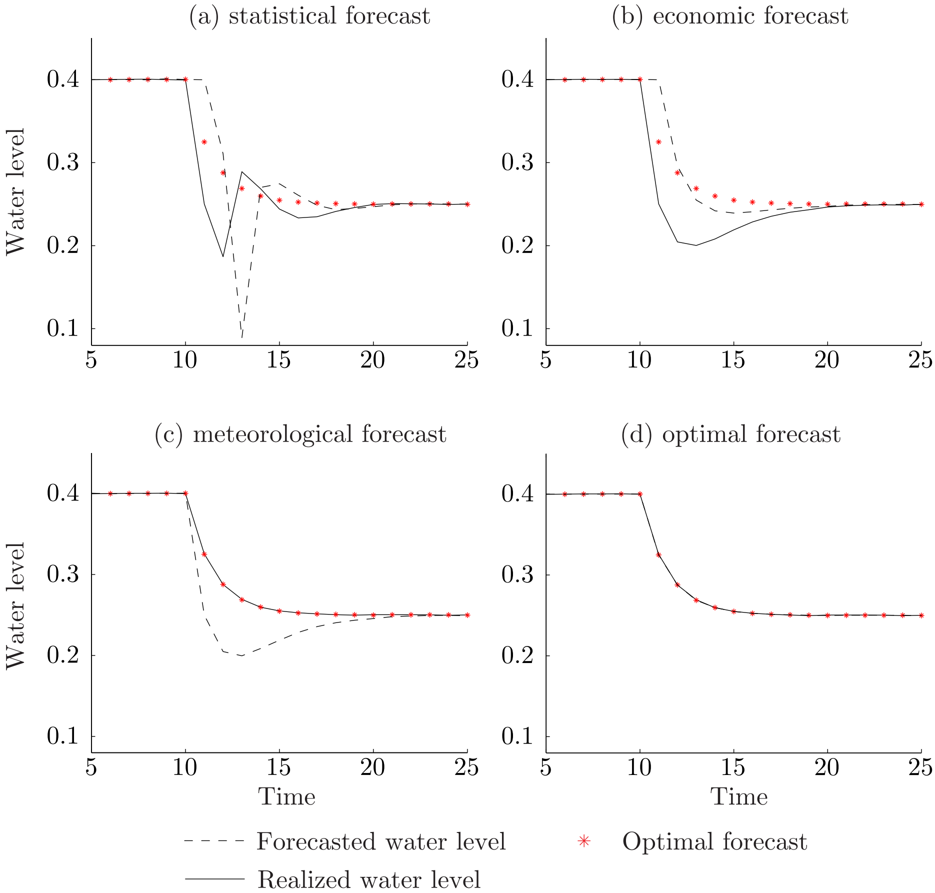

Fig. 1. Forecast performance during drought

exceeds the water supply, thus causing the water level to drop. In subsequent periods, the forecaster updates his or her statistical model, incorporating this sudden drop in water. Updating leads to a forecast of extremely low levels of water. This forecast now turns out to be too pessimistic, since the low forecast will cause most farmers to shut down, decreasing the demand for water. This type of behavior recurs multiple times before the system finally settles down about 10 periods after the shock.

This case illustrates the large forecast errors that can occur if structural breaks occur and agents are naïve about supply, but it is perhaps unreasonable to assume that farmers do not know the structure of the water system. Panel (b) illustrates the outcome when farmers are fully rational and sophisticated but still experience an unknown structural break. Without knowing that the drought will hit, the forecaster and farmers initially believe in an overly-optimistic forecast, causing excess demand and a deep fall in the water level. Due to the exponential kernel, older rainfall observations are downweights, so eventually the pre-break rainfall data is no longer used for estimation, and forecasts match realizations. Since farmers and the forecaster know demand, the forecast errors in this case are due to the break in the rainfall process rather than feedback from expectations.

In contrast, the meteorological forecaster in panel (c) experiences forecast errors that are purely due to feedback from expectations. Knowing the rainfall process, the forecaster immediately revises the forecast down in expectation of the drought arriving. This immediate revision, however, does not take into account the fact that a lower forecast for rainfall will reduce demand. Indeed, the low rainfall forecast causes such a broad shutdown of farms that the overall water level declines by less than the forecast. This dynamic continues for multiple periods while the forecaster incorporates the new demand data. One can see that even perfect knowledge of the physical system can, in such a case, lead to a high degree of bias.

We have focused on a change in the rainfall process because comparison of the meteorological and economic forecasts in this case offers a clear illustration of the errors caused by expectation feedback. We could also consider changes in other exogenous model parameters. One of particular interest is an increase in price, which would motivate more farmers to enter the market, increasing total demand for water. In this case, the meteorological forecaster will again experience forecast errors due to feedback, but in the opposite direction.

Price changes also highlight the role of a key simplification in the current model—that the price of the agricultural good is exogenous to the quantity of the agricultural good supplied. If the farmers being modeled represent a small enough part of the overall market for the agricultural good, this assumption is plausible since these farmers have a "negligible" effect on the overall supply and hence price of the good. If this assumption is deemed implausible, we would need to estimate the demand curve, which relates the price and quantity of the agricultural good. Different predictions can emerge depending on the slope of the demand curve. If a small decrease in quantity of the agricultural good leads to a large increase in price, a drought can theoretically cause more farmers to plant crops during a drought. The intuition for this result is that expected profits for a farmer depend both on the probability of a successful harvest and the price when the harvest is successful. The drought makes it less likely for the harvest to be successful, but a farmer might reason that the price will be higher during a drought because *other* farmers' crops may fail, so total supply of the good will fall and the price will rise. If the price is high enough, a farmer might then *intensify* his planting in the presence of a drought.

Empirically, price changes and weather shocks can combine to produce unexpected outcomes. For instance, in the mid-1900s, technological advancement led to cheaper, more widely available ground water in the Midwest of the United States. This change increased the available supply of water, reduced the price of producing agricultural goods, and potentially provided a back-stop supply of water in case of prolonged drought. Farmers reacted to these changes by adding new farmland and using less drought-resistant crops, leading to a zero net change in the ability of farmers to cope with drought [9].

2.5 Comparing forecast quality

Figure 1 implicitly compares each forecast to the optimal structural forecast. We can also assess forecast quality more formally from a statistical perspective. The main point of our analysis is that even with sophisticated agents and forecasters who accurately forecast the underlying physical system, feedback can lead to forecast errors. Figure 2 panel (a) summarize this error by calculating the squared error each period for the statistical, economic, and meteorological forecasts and then taking the

Panel (a) shows period-by-period RMSE for each of the non-structural forecasts. Panel (b) shows the cumulative value of the RMSE in the first panel, starting at the date of drought onset.

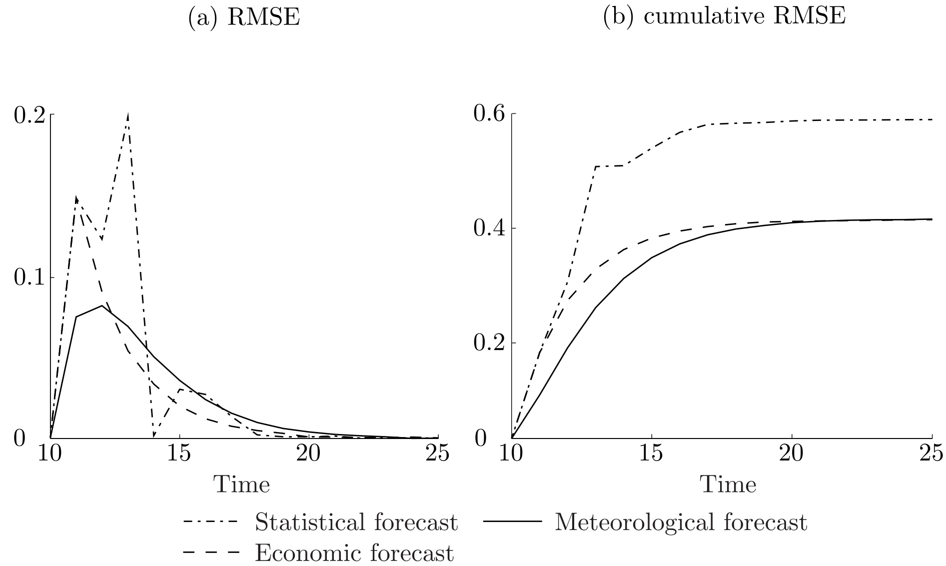

Fig. 2. Statistical and economic performance of the forecasts

square root of the average of the squared error across simulation draws. The resulting value can be interpreted as a root mean squared error of each forecast.

One can see that during the initial period after drought onset, both the economic and statistical forecaster perform relatively worse than the meteorologist. After one period of high error, the economic forecaster does progressively better, but the statistical forecaster continues to experience high and variable forecast performance. By construction of our exponential kernals, 10 to 20 periods after the shock, all forecasters cease to experience substantial errors.

In terms of cumulative error, Figure 2 panel (b) shows that in the long run, the economic and meteorological forecasters perform equally well. Since the economic forecaster only experiences a structural shock, this equivalence implies that forecast feedback can be as bad for forecast errors as the much more widely studied topic of structural breaks. In the presence of multiple forecasts of potentially differing quality, a modelling approach as in [7] could potentially be taken. For simplicity we consider forecasts and water level dynamics under differing forecasting regimes rather than with competing forecasters.

The non-structural forecasters in our model are each constrained in ways that lead to forecast errors. How reasonable are these constraints? In practice, many forecasts exhibit persistent bias even in cases where expectation feedbacks are not present and where there is some understanding of how to predict the underlying physical system. A previously mentioned example is forecasts of El Niño/Southern Oscillation (ENSO). Although the first skillful ENSO forecast was published in 1986, forecast bias was still present 14 years later [2,4]. Daily precipitation forecasts issued by local meteorologists have also been shown to exhibit substantial bias toward predicting more rainfall than actually realized. Notably, some of the same bias is also present in national forecasts made by the Weather Channel, a fact documented by researchers and popularized by Nate Silver [1,18]. The Weather Channel has the largest market share of any consumer forecasting service in the United States [10]. Finally, in the context of the model, forecasts issued by NOAA failed to anticipate the onset of the current drought in California and the western United States.[4] Lack of predictability of onset is common to multiple drought forecasting models [21].

Other cases exhibit both forecast bias and the potential for feedback. A well studied example comes from electricity supply forecasts issued by the North American Electric Reliability Council (NERC). From 1974 to 1981, the NERC issued projections of national electricity supply that were invariably too high [12]. The Federal Reserve's Greenbook forecasts of economic growth have also been shown to exhibit consistent bias [13]. In an example combining both physical systems and feedbacks similar to those used in the model, rapid warming in the New England cod fishery combined with a failure to correctly predict fish stock led to consistently over-optimistic stock forecasts and the collapse of the fishery [14]. It is unclear whether forecast bias in each of these cases is due to expectation feedback, an objective function of the forecaster that trades bias for other forecast characteristics, or some other factor, but differentiating between these different sources of error would be a potentially fruitful empirical exercise.

Finally, it is important to note that forecast error and economic loss need not be the same or even similar. In our simple model, even though the meteorological forecast is systematically biased downward as shown in Figure 1 panel (c), the path of realized water level is exactly that of the optimal forecast. In the model this occurs because we assume the agents are sophisticated enough to predict demand. Thus, the agents combine their own information with the forecast level of rainfall to attain all the information they need. The forecast of water level is systematically biased, but this comes at no cost to the agents relative to the optimal forecast. In practice, demand depends on many factors outside the purview of agents and thus will be forecasted imperfectly. In that case, agents might experience economic losses even with perfect meteorologic forecasts.

3 Conclusion

Climate change is expected to lead to more frequent and more intense weather extremes. Physical scientists are hard at work to improve our ability to forecast these events. Social scientists are similarly engaged in efforts to understand how individuals will react to these extremes. How individuals will respond to ever-improving forecasts, and what that implies about the ultimate accuracy and impacts of those forecasts, is an open question.

In this paper, we have developed a highly simplified model of an agricultural setting to illustrate what happens when the actions of economic agents respond to forecasts about physical systems and how that, in turn, influences the accuracy of model predictions under a variety of forecasting scenarios. The key insight from our analysis—that feedback between forecasts and agent expectations can render forecast invalid *ex post*—should extend to a wide range of real world settings, including those outside the realm of climate extremes.

For example, forward guidance on interest rates in financial markets shapes saving and investment decisions that can impact equilibrium costs of borrowing. Similarly, biological stock assessments can impact resource extraction and thus the population dynamics of the species of interest. Even forecasts at small temporal scales, such as that of electricity production can impact price expectations and thus capital investments and long-term energy demand by consumers. Assessing the degree to which expectations are incorporated in forecasts in these sectors and improvements to forecasts that do not incorporate expectations offer rich areas for future research. Greater

[4] See the forecast verification from the onset of the drought here: `http://www.cpc.ncep.noaa.gov/products/expert_assessment/sdo_verification/2011/NDJ11-initial.png`

collaboration between social and physical scientists is needed to deepen our understanding of forecast consumption, expectation formation, and the ways in which those expectations may alter the dynamics of systems critical to societal health and well being.

Shrader gratefully acknowledges support from California Sea Grant College Program, Project No. E/MRE-8 and the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1144086.

References

1. J. Eric Bickel and Seong Dae Kim. Verification of The Weather Channel Probability of Precipitation Forecasts. *Monthly Weather Review*, 136(12):4867–4881, dec 2008.
2. MA Cane, SE Zebiak, and SC Dolan. Experimental forecasts of EL Ni\~no. *Nature*, 1986.
3. Christophe Chamley. *Rational herds: Economic models of social learning*. Cambridge University Press, 2004.
4. Dake Chen, MA Cane, SE Zebiak, Rafael Canizares, and Alexey Kaplan. Bias correction of an ocean-atmosphere coupled model. *Geophysical Research Letters*, 27(16):2585–2588, 2000.
5. B. I. Cook, T. R. Ault, and J. E. Smerdon. Unprecedented 21st century drought risk in the American Southwest and Central Plains. *Science Advances*, 1(e1400082):1–7, feb 2015.
6. Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.
7. Matthew Gentzkow and Emir Kamenica. Competition in Persuasion. 2015.
8. James D Hamilton. *Time Series Analysis*. Princeton University Press, 1 edition, 1994.
9. Richard Hornbeck and Pinar Keskin. The Historically Evolving Impact of the Ogallala Aquifer: Agricultural Adaptation to Groundwater and Drought. *American Economic Journal: Applied Economics*, 6(1):190–219, jan 2014.
10. IBISWorld. Clear skies : Niche services and mobile apps are forecast to drive industry demand IBISWorld Industry Report OD5752 Weather Forecasting Services in the US About this Industry. Technical Report October, 2012.
11. Emir Kamenica and Matthew Gentzkow. Bayesian Persuasion. *American Economic Review*, 101(6):2590–2615, oct 2011.
12. Charles R. Nelson and Stephen C. Peck. The NERC Fan: A Retrospective Analysis of the NERC Summary Forecasts. *Journal of Business & Economic Statistics*, 3(3):179–187, jul 1985.
13. Andrew J. Patton and Allan Timmermann. Forecast Rationality Tests Based on Multi-Horizon Bounds. *Journal of Business & Economic Statistics*, 30(1):1–17, jan 2012.
14. Andrew J. Pershing, Michael A. Alexander, Christina M. Hernandez, Lisa A. Kerr, Arnault Le Bris, Katherine E. Mills, Janet A. Nye, Nicholas R. Record, Hillary A. Scannell, James D. Scott, Graham D. Sherwood, and Andrew C. Thomas. Slow adaptation in the face of rapid warming leads to collapse of the Gulf of Maine cod fishery. *Science*, page aac9819, oct 2015.
15. Mark Rosenzweig and Christopher R Udry. Forecasting Profitability. 2013.
16. Suranjana Saha, Shrinivas Moorthi, Xingren Wu, Jiande Wang, Sudhir Nadiga, Patrick Tripp, David Behringer, Yu-Tai Hou, Hui-ya Chuang, Mark Iredell, et al. The ncep climate forecast system version 2. *Journal of Climate*, 27(6):2185–2208, 2014.
17. Jeffrey Shrader. Estimating adaptation using forecasts. *Working paper*, 2015.
18. Nate Silver. *The Signal and the Noise: Why so many predictions fail - but some don't*. Penguin, New York, NY, USA, 1 edition, 2012.
19. Allan Timmermann and Clive W. J. Granger. Efficient market hypothesis and forecasting. *International Journal of Forecasting*, 20(1):15–27, jan 2004.

20. United States Department of Agriculture. Irrigation & Water Use, 2015.
21. Xing Yuan and Eric F. Wood. Multimodel seasonal forecasting of global drought onset. *Geophysical Research Letters*, 40(18):4900–4905, sep 2013.