LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

UCRL-JRNL-216578

Weather-based forecasts of California crop yields

D.B. Lobell, K.N. Cahill, C.B. Field

October 27, 2005

California Agriculture

Weather-based forecasts of California crop yields

David B. Lobell[1*], Kimberly Nicholas Cahill[2,3], and Christopher B. Field[2]

Submitted to California Agriculture. Please do not distribute or cite without author permission

[1] Energy and Environment Directorate, Lawrence Livermore National Laboratory, Livermore, CA 94550

[2] Department of Global Ecology, Carnegie Institution, Stanford, CA 94305

[3] Interdisciplinary Graduate Program in Environment and Resources, Stanford University, Stanford, CA 94305

* Author for correspondence and reprints: David B. Lobell
Lawrence Livermore National Laboratory
P.O. Box 808 L-103
Livermore, CA 94550
Tel: (925) 422-4148
Fax: (925) 423-4908
Email: dlobell@llnl.gov

Abstract

Crop yield forecasts provide useful information to a range of users. Yields for several crops in California are currently forecast based on field surveys and farmer interviews, while for many crops official forecasts do not exist. As broad-scale crop yields are largely dependent on weather, measurements from existing meteorological stations have the potential to provide a reliable, timely, and cost-effective means to anticipate crop yields. We developed weather-based models of state-wide yields for 12 major California crops (wine grapes, lettuce, almonds, strawberries, table grapes, hay, oranges, cotton, tomatoes, walnuts, avocados, and pistachios), and tested their accuracy using cross-validation over the 1980-2003 period. Many crops were forecast with high accuracy, as judged by the percent of yield variation explained by the forecast, the number of yields with correctly predicted direction of yield change, or the number of yields with correctly predicted extreme yields. The most successfully modeled crop was almonds, with 81% of yield variance captured by the forecast. Predictions for most crops relied on weather measurements well before harvest time, allowing for lead times that were longer than existing procedures in many cases.

Introduction

Forecasts of crop yields can provide important information about commodity markets, and are frequently used by growers, industry, and governments to make decisions (Vogel and Bange, 1999). Growers may use forecasts, for instance, to plan for their harvest, storage, and distribution strategies. California growers used the 2004 forecast of a large rice harvest to arrange greater storage capacity, and may be using the smaller 2005 almond forecast to allocate limited quotas among preferred customers (D. Flohr, CASS, pers comm). Similarly, industries involved in the handling and trading of commodities often use information on future harvests to make various logistical decisions (Hammer et al., 2001).

Each year, the California Agriculture Statistics Service (CASS) estimates the size of the coming harvest for various major California crops, including almonds, grapes, olives, oranges, and walnuts (NASS, 2005a; NASS, 2005b). These estimates are categorized as either subjective or objective. The former are based on phone interviews with hundreds of farmers to assess their opinion of crop development, while the latter are based on field samples taken from hundreds of fields. Forecasts are generally made public one to three months before the end of harvest (see the California Objective Measurement Report, http://www.nass.usda.gov/ca/rpts/om/indexom.htm, for more information).

It is well known that one of the main factors causing yields to change from year to year is climate variability – no two growing seasons experience exactly the same weather. Indeed, grower expectations of crop yields are likely based at least partially on subjective weather observations and perceived relationships between weather and yields. To our knowledge, however, there currently is not any objective, quantitative use of weather measurements in existing yield forecast procedures. Such an approach would be attractive because yields could potentially be forecast with lower costs, higher accuracies, and/or longer lead times.

Building Forecast Models

To test the ability of weather measurements to forecast crop yields prior to harvest, we studied the statistical relationships between historical weather and crop yield records. We selected 12 crops (wine grapes, lettuce, almonds, strawberries, table grapes, hay, oranges, cotton, tomatoes, walnuts, avocados, and pistachios) that are among the most valuable crops in California (Table 1; California Agricultural Statistics Service, 2004a), and obtained state yield data for 1980-2003 from the California County Agricultural Commissioners (California Agricultural Statistics Service, 2004b). Since several crops have exhibited significant positive yield trends since 1980 owing to management and crop cultivar changes, we removed a linear trend from each crop to produce a time series of yield anomalies, or departures from expected yields. A positive anomaly indicates yields higher than expected based on time trend, while a negative anomaly indicates below- expected yields.

Daily weather records for the same time period were obtained for 382 stations throughout California from the California Climate Change Center at the Scripps Institution of Oceanography (M. Tyree, pers comm.). The average daily minimum and maximum temperature and precipitation for each month in each county was then computed, resulting in a monthly time series of 3 variables for 24 years. For each crop, a state-wide monthly time series for each meteorological variable was then calculated by weighting each county by the relative area of the crop in that county in 2003 (for more details, see Lobell et al., in review).

The weather and yield data were then combined in linear regression models to test how well yield anomalies could be predicted before harvest based on monthly weather measurements. Between two and four weather variables were selected for each crop based on a combination of objective (good model accuracy) and subjective (physiologically reasonable) criteria. Since temperature and precipitation can have a non-linear effect on yields, with yields maximized at intermediate values, we included the squared values of the weather variables in the regression model along with the variables themselves. The selected variables and months for each crop are given in Table 2. For crops such as pistachios known to exhibit alternate bearing, with years of high reproductive growth (high yields) alternating with years of high vegetative growth (low yields), yield anomalies from previous years were also included in the model. The total number of predictors, including the weather variables, squared variables, and previous years' yields, ranged from four to eight. (The model equations are omitted for brevity, but can be obtained by contacting the author.)

Comment [KNC1]: I think "monotonic" is too technical a term, even though we define its meaning in the next phrase.

An important step when building statistical models is to independently test model predictions, since tests using the same data used to calibrate the model will tend to be overly optimistic (Hastie et al., 2001). The straightforward approach of reserving part of the data during model calibration, however, is problematic when data quantity is limited. An alternative approach that we employ here is leave-one-out cross-validation, whereby a single year is "left out" of the calibration step and then subsequently compared to model prediction in that year. This comparison is done for each year, in this case resulting in 24 comparisons between model predictions and observations.

Forecast Accuracy

The results of the cross-validation analysis suggest that yields of some crops can be forecast with fairly high accuracy based on objective weather measurements (Figure 1 and Table 3). For many crops, close to or greater than 50% of the variability in yield anomalies was captured by the model forecasts, meaning that the selected weather variables explained over half of the variations observed in crop yields. Interestingly, the models did fairly well at forecasting extremely low yields, such as almonds in 1995, oranges in 1991, and tomatoes and cotton in 1998 (Figure 1). Almonds were particularly well modeled, with over 80% of variance captured by the model.

Comment [KNC2]:

For a few crops, some of the power of the models came from knowing the previous year's yield (Table 3). For instance, including weather information did not improve the pistachio model, where the biological pattern of alternate bearing seemed to dominate effects on yield more than any weather signal. For all other crops, however, most or all of the predictive power came from weather variables.

As an alternative measure of forecast skill, we considered the fraction of years in which the model forecast was closer to the yield anomaly than zero (Table 3). That is, we examined the frequency with which the model correctly predicted whether the yield would be above or below the trend. For a random forecast, this statistic has a distribution whose mean is 0.5 and whose 90^{th} percentile is 0.625 for a 24-year record (15 out of 24 years). Thus, six of 12 crops had a forecast with skill greater than a random forecast using

Comment [KNC3]: This was a fragment w/o a verb before.

this criterion and significance level, with three others (wine grapes, table grapes, and lettuce) slightly below this value.

Another criterion is the ability of forecasts to correctly predict unusually high or low yields, which may be of particular interest to some forecast users. For each year, both the forecast and the actual yield were classified into one of 4 classes: below 1 standard deviation (SD) from zero, between -1 SD and zero, between zero and 1 SD, and above 1 SD from zero. The first and fourth of these classes represent unusually low or high yields, respectively, while the middle two represent more moderately negative or positive years. We then computed the number of years when the forecast correctly predicted the yield class, was off by one class (in either direction), two classes, or three.

Comment [KNC4]: I think acronyms should be in caps, no?

Most crops did not exhibit any years when the forecast was off by more than one class. There were some exceptions, for example lettuce yields in 1981 were forecast to be slightly negative but were actually very high (above 1 SD), and the reverse was true for hay in 1995. Overall, however, it appears that the forecasts were usually no more than 1 class off, meaning that most of the cases above where the forecast predicted an anomaly in the wrong direction corresponded to years with moderate yields, so the forecast was in fact not far from the observed yield. None of the crops exhibited any years with a forecast error of three classes.

To test the significance of these class accuracies, we performed 1000 simulations using two normally distributed random variables of length 24 years. The average percent of years with an error of 0, 1, 2, or 3 classes was 28%, 45%, 22%, and 5%, respectively. Only 10% of the simulations had greater than 40% of years (10 out of 24) classified correctly by chance, while all crops except strawberries, pistachios, and walnuts met this criteria. This indicates that the forecast accuracies for most of the crops were statistically significant by this measure.

Forecast Timing

Forecast timing can be as important as its accuracy. A “forecast” made after harvest, for example, would not be very valuable. As shown in Table 3, most of the models are capable of providing forecasts at least several months before the end of harvest, giving growers or others an opportunity to use forecast information in making their decisions. For instance, the models for almonds and walnuts relied mainly on winter weather, while harvest does not begin until late summer.

A comparison between the months when these models can forecast yields with the months that currently available crop yield forecasts are released by the USDA is shown in Table 3. The two approaches were similar for wine grapes, table grapes and cotton, and existing forecasts were available four months earlier for processing tomatoes than our models provided. However, our models offer significant timing advantages over existing forecasts for almonds (3-4 months earlier than current forecasts), hay (2 months earlier), strawberries (5 months earlier), and walnuts (7 months earlier).

Potential Improvements

The current analysis was limited to only a dozen of the many crops grown in California, and considered only state-wide yields. For several crops we have also chosen to aggregate over different sub-crop groupings, such as in combining varieties of hay and lumping navel and Valencia oranges. In addition, we have utilized only monthly averages of three meteorological variables (number of frost days per month was also considered, but did not substantially improve any of the models).

These decisions reflect an explicit desire to test forecasts of state yields for major crops using commonly reported climatic data. However, data for many additional crops are currently available at both state and county levels, as are additional weather measurements at time scales from hourly to monthly. Thus, an open question is how well other crops can be modeled and whether different scales of analysis and meteorological indices would substantially improve forecast accuracies. Additional information such as remote sensing data might also aid predictions.

Comment [KNC5]: "Ancillary" sounded too technical

It is also possible that different model formulations could improve results. For example, in certain situations process-based models that rely on mechanistic understanding of crop growth and yield may out-perform statistical models such as the ones developed here, which are derived from observed relationships without explaining the mechanism causing the relationship. Alternative statistical approaches to the multiple linear regression approach used here may also improve accuracies. (For example, we tested the use of regression trees, which did not perform as well.) Whether these more sophisticated approaches offer worthwhile improvements can only be tested on a case-by-case basis, using actual observations and well-defined criteria for an ideal forecast.

Conclusions

The models developed in this study demonstrate promise for forecasting statewide crop yields based on weather measurements. As the significance level of the models depended on the specified performance criteria, it is clear that the eventual utility of such forecasts will depend on the acceptable types and magnitude of errors for a particular application. The potential to forecast yields also exhibits an obvious dependence on crop type. In general, almonds exhibited significantly greater forecast accuracies than the other crops considered here. As California almonds represent the most valuable export crop in the state and comprise over 80% of global almond production (Almond Board of California, 2004), such forecasts could be of great relevance to almond trade and management decisions. For example, an almond grower could have used data on January rainfall and February nighttime temperatures to correctly predict the low yield in early March of 1995 and adjust cultural or marketing practices accordingly, well before the forecasts from USDA become available in May and June in advance of fall harvests.

While field-based surveys are likely to be more accurate than weather-based forecasts, it is important to consider the tradeoff between forecast accuracy, cost, and timing. The low cost and long lead-times possible with weather-based models would likely provide a useful complement to existing approaches for crops that are currently surveyed. For crops that are not currently forecast by USDA, such as avocados, these models present an opportunity to develop forecasts with minimal cost by using existing weather measurements.

Acknowledgements

The authors are grateful to Mary Tyree for providing the climate data and Doug Flohr for information on current forecast applications. Funding for this work came in part from an Environmental Protection Agency (EPA) Science to Achieve Results (STAR) graduate fellowship to DL, a David and Lucile Packard Foundation Stanford Graduate

Fellowship to KNC, and the Office of Science, Biological and Environmental Research Program (BER), US Department of Energy (DOE), through the Western Regional Center of the National Institute for Global Environmental Change (NIGEC). This work was performed under the auspices of the U.S. Department of Energy, National Nuclear Security Administration by the University of California, Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

References

Almond Board of California. 2004. Almond Almanac. Almond Almanac. Almond Board of California, Modesto, CA.

California Agricultural Statistics Service. 2004a. California Agricultural Statistics 2003: Overview. California Agricultural Statistics 2003: Overview. USDA-NASS, Sacramento.

California Agricultural Statistics Service. 2004b. California Agricultural Commissioners' Data [Online]. California Agricultural Commissioners' Data. http://www.nass.usda.gov/ca/bul/agcom/indexcac.htm.

Hammer, G.L., J.W. Hansen, J.G. Phillips, J.W. Mjelde, H. Hill, A. Love, and A. Potgieter. 2001. Advances in application of climate prediction in agriculture. Agricultural Systems 70:515-553.

Hastie, T., R. Tibshirani, and J. Friedman. 2001. The elements of statistical learning: data mining, inference, and prediction Springer, New York.

Lobell, D.B., K.N. Cahill, and C. Field. In review. Historical effects of temperature and precipitation on California crop yields. Climatic Change.

NASS. 2005a. California Subjective Measurement Reports, 1997-2005 [Online]. California Subjective Measurement Reports, 1997-2005. Available by California Field Office, National Agriculture Statistics Service, United States Department of Agriculture http://www.nass.usda.gov/ca/oth/indexfor.htm.

NASS. 2005b. California Objective Measurement Reports, 1997-2005 [Online]. California Objective Measurement Reports, 1997-2005. Available by California Field Office, National Agriculture Statistics Service, United States Department of Agriculture http://www.nass.usda.gov/ca/rpts/om/indexom.htm.

Vogel, F., and G. Bange. 1999. Understanding crop statistics [Online]. Office of the Chief Economist, U.S. Department of Agriculture, Miscellaneous Publication No. 1554. http://www.usda.gov/nass/nassinfo/pub1554.htm.

Figure Legends

Figure 1. Observed (•) and forecasted (x) yields for 1980-2003. Forecasts are derived for each year using a model fit to data in all other years. Coefficient of determination (R^2) between observed and forecasted yields is shown for each crop.

Table 1. Economic value and national importance of production of crops studied.

Crop	2003 Value ($ Millions) [a]	% of U.S. Production
Grapes, Wine	1828	96%
Lettuce	1634	88%
Almonds	1506	99%
Strawberries	973	83%
Grapes, Table	953	91%
Hay	950	12%
Oranges	949	22%
Cotton	774	10%
Tomatoes, Processing	571	95%
Walnuts	434	99%
Avocados	402	95%
Pistachios	173	99%

[a] Values are taken from CASS (2004b), which are based on free-on-board (FOB) prices that include value added items such as packing and inspections.

Table 2. Months and weather variables used for yield forecasts for each crop evaluated in this study.

	Year prior to harvest					Year of harvest								
Crop	Aug	Sept	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sept
Grapes, wine		ppt							tmn		ppt			
Lettuce			tmx				tmx		tmx					
Almonds						ppt	tmn							
Strawberries				all										
Grapes, table			ppt			ppt			tmn			tmn		
Hay							ppt				ppt			
Oranges					tmn					ppt				
Cotton										tmx	tmn			
Tomatoes									tmx		tmx			tmn
Walnuts				tmx			ppt							
Avocados	tmx		ppt							tmn				

tmn = average minimum temperature; tmx = average maximum temperature; ppt = total rainfall; all = all three variables

Table 3. Summary of forecast accuracy and timing for crops evaluated in this study.

Crop	R^2cv	RMS (%)	R^2 using only previous yields [a]	Fraction of years with forecast in correct direction	Last month used in forecast	Month of USDA forecast[b]	Peak harvest period	# months between forecast and end of harvest
Grapes, Wine	0.59	6.4	*	0.61	June	July-August	August - October	4
Lettuce	0.44	4	*	0.61	April		continuous	--
Almonds	0.81	7.8	0.17	0.73	February	May (subjective); June (objective)	August-October	8
Strawberries	0.49	4.6	*	0.48	previous November	April	Continuous	--
Grapes, Table	0.62	6.7	*	0.61	July	July-August	July-September	4
Hay	0.44	3.9	0.01	0.55	June	August	March-November	5
Oranges	0.69	8.8	0.22	0.68	May	Navel: September; Valencia: March	November-May, May-Oct[c]	6
Cotton	0.56	6.3	*	0.54	June	June-August	October-December	6
Tomatoes, Processing	0.49	3.1	*	0.67	September	May & September	June-November	2
Walnuts	0.43	7.3	0.06	0.57	February	September	September-November	9
Avocados	0.57	16.7	*	0.7	May		Continuous	--
Pistachios	0.35	27.5	0.42	0.7	n/a	August	September-November	--

R^2cv = cross-validated R^2 , the proportion of yield variance explained by the weather predictor variables

RMS = root mean squared difference between forecast and observed yield, expressed as a percentage of average yield for 2000-2003

[a] Only crops that exhibited alternate bearing were modeled with previous years' yields.

[b] Available in California Crop Production Reports (at http://www.nass.usda.gov/ca/)

[b] The first period refers to Navel orange harvest, and the second to Valencia oranges.

Formatted: Not Highlight

Comment [KNC6]: Maybe offer a phrase about how to use this to interpret accuracy of forecasts.

Figure 1

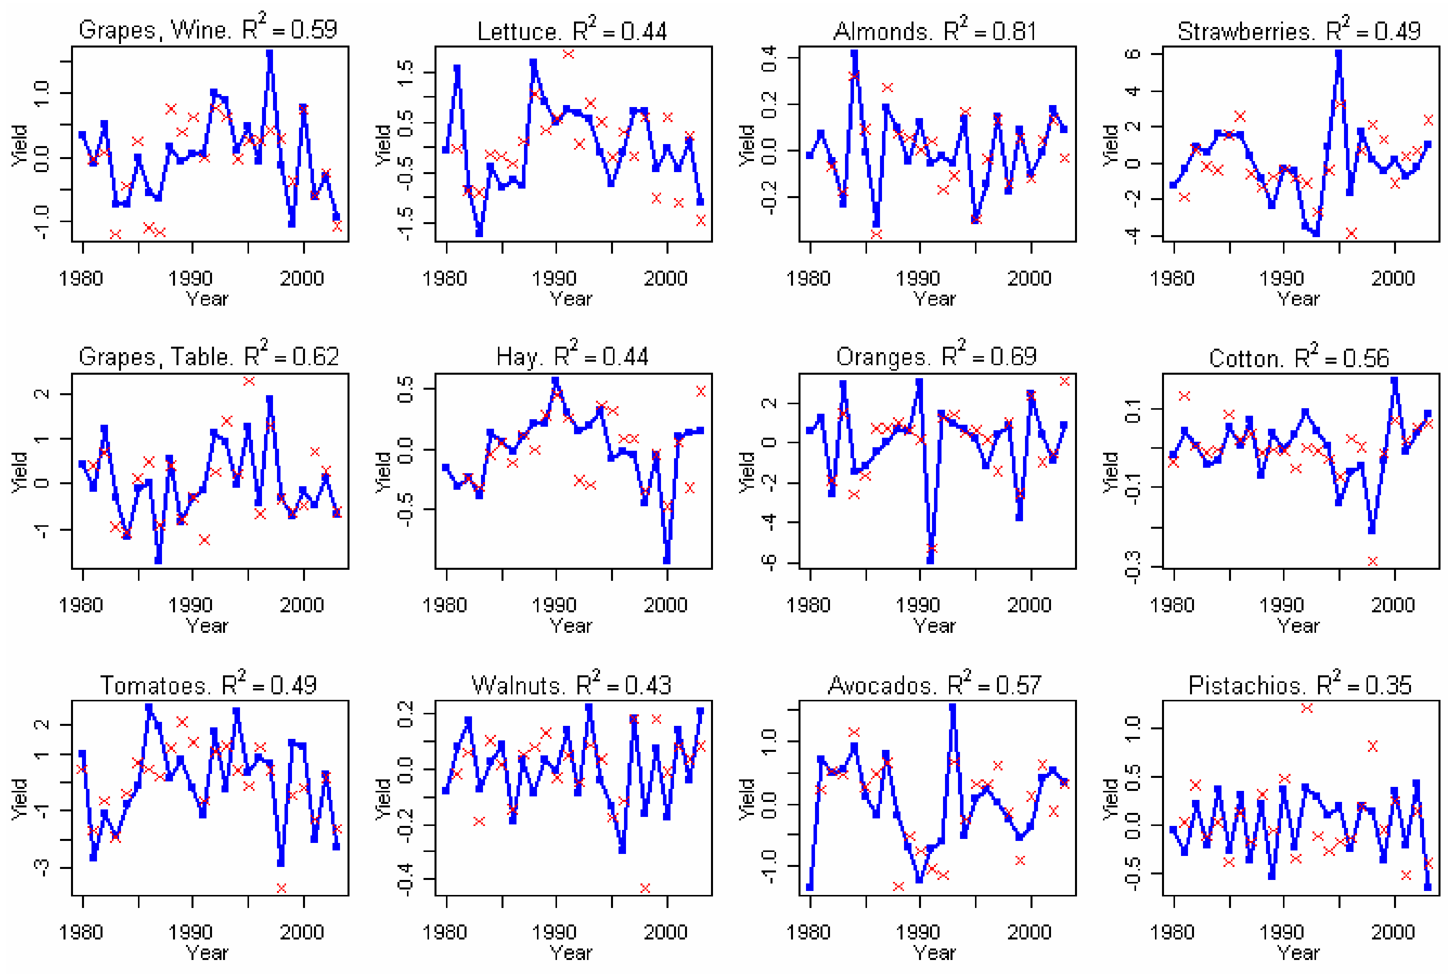

Grapes, Wine. $R^2 = 0.59$
Lettuce. $R^2 = 0.44$
Almonds. $R^2 = 0.81$
Strawberries. $R^2 = 0.49$
Grapes, Table. $R^2 = 0.62$
Hay. $R^2 = 0.44$
Oranges. $R^2 = 0.69$
Cotton. $R^2 = 0.56$
Tomatoes. $R^2 = 0.49$
Walnuts. $R^2 = 0.43$
Avocados. $R^2 = 0.57$
Pistachios. $R^2 = 0.35$
Yield
Year
1980
1990
2000