35th Annual conference of South African Society for Atmospheric Sciences
SASAS2019
http://www.arc.agric.za/arc-iscw/SASASconference/Pages/default.aspx

How costly are poor seasonal forecasts?

Willem A. Landman, Department of Geography, Geoinformatics and Meteorology, University of Pretoria

Emma Archer, Department of Geography, Geoinformatics and Meteorology, University of Pretoria

Mark Tadross, Climate Systems Analysis Group, University of Cape Town

Abstract

Seasonal forecast system development has made significant advances in recent years, including the development of models for hydrological, agricultural and health applications utilising forecast output from complex global climate models. The skill levels of these models are in some cases (location and season) found to be promising when evaluated over an extended period of time – notwithstanding the fact that during some years forecasts may still be wrong even for skilful models. In this study, we investigate what the financial implications might be when forecast crop-yields are wrong in consecutive years. The paper first introduces a linear statistical dry-land crop-yield model that uses output from a coupled ocean-atmosphere climate model as a predictor of crop-yield. The crop model is shown to be skilful, but produced poor forecasts for three consecutive years during the test period. We evaluate what the possible cost implication might be for a farmer who makes financial investments or takes financial risks in proportion to the crop-yield forecasts.

Keywords: crop-yield forecasts, coupled ocean-atmosphere model, El Niño and La Niña, accumulated profit

Introduction

The seasonal forecast community has developed complex climate models for operational seasonal forecasting in South Africa (Beraki et al., 2014). For optimal seasonal forecast production, atmospheric models are coupled to similar models for the ocean, the land surface and sea ice. Notwithstanding their demonstrated accuracies, statistical correction methods are recommended even for today's coupled climate model forecasts (Barnston and Tippet, 2017). The use of such multi-tiered forecast systems have shown to be more accurate for seasonal rainfall forecasts, at least for SADC (Landman et al., 2012). Recently, hindcasts (or re-forecasts) over a period spanning several decades have been used in the development of application models for agriculture in southern Africa (Malherbe et al., 2014) and hydrology (Muchuru et al., 2014). In some cases, the developed application models have been used in an operational seasonal forecasting environment (see the archived forecasts produced by the University of Pretoria for examples: https://tinyurl.com/ybrb3a72). As is the practice with operational seasonal forecasting, these forecasts, including applications forecasts, are accompanied by some indication of forecast skill evaluated over an independent test period. The skill estimates represent a general statement on the overall skill of the forecast system. In this study, we want to develop an application model, and more specifically a model for the prediction of dry-land crop-yields at a single farm in South Africa. We then determine the skill levels of the model, followed by an assessment of possible financial implications for the forecast user (in the agricultural sector) when there is a succession of poor or "missed" forecasts produced by the model.

Data and Methodology

a. Data

A set of coupled model hindcasts (or re-forecasts) and end of season crop-yield data are used in the following analysis. The climate model data have been used already for a number of predictability studies (e.g. Landman et al, 2012) and consists of ensemble mean (from 12 members) 850 hPa geopotential height (i.e. near-surface atmospheric circulation) as a proxy for rainfall hindcasts. This geopotential height field has long since been established as a predictor that can replace a climate model's predicted rainfall fields in statistical downscaling for southern Africa (Landman and Goddard, 2002). The geopotential height anomalies are forecast using the ECHAM4.5-MOM3 coupled model (DeWitt, 2005) for December to February (DJF) seasons, with a model initialization month of November. Since DJF is often the best forecast skill season over the region and the rainfall during this season plays a significant role during grain filling and tasselling, it was decided to use only climate model data for DJF as a predictor of end of season dryland crop-yield. Even though additional atmospheric variables (e.g. relative or specific humidity) might improve forecast skill, we leave aside additional variables as our focus is on demonstrating the impact of poor forecasts, and geopotential height is sufficient for deriving a reasonably skilful model.

The crop-yield data have been obtained from a farm near Bapsfontein (26°0′0″S 28°26′0″E) in South Africa. The period over which the analysis is done is the 21 years from 1987/88 to 2007/08. This period consists of 7 El Niño, 7 La Niña and 7 ENSO-neutral seasons according to the Oceanic Niño Index for cold

Corresponding author: Tel: +27-12-420-3713; Email: Willem.Landman@up.ac.za

and warm episodes. The crop yields are first detrended to remove the linear upward trend often associated with crop-yield data. To ensure that the yield data are from a normal distribution for optimal statistical modelling to be performed, the natural logarithm of the yield values are calculated and referred to in the following text. The Lilliefors (Wilks, 2011) goodness-of-fit test shows that the newly derived crop yield values are indeed from an unspecified normal distribution.

b. Methods

The IRI's Climate Predictability Tool (CPT) is used for producing crop-yield hindcasts and for verification. The predictor set is the ensemble mean 850 hPa geopotential height field of the coupled model, and the predictand the Bapsfontein crop yields. The canonical modes of the hindcasts are used in a multiple linear regression model as predictors. The forecast skill level of the statistical crop-yield model is first tested using a cross-validation setup with a 5-year-out window. Then the crop model is used to produce retro-active crop-yield forecasts for the two periods from 1999 to 2008 and from 2003 to 2008. The retro-active forecast process of the CPT creates probabilistic forecasts over these periods for three equi-probable categories with thresholds defined by respectively the 33.3rd (below which is the low yield category) and 66.7th (above which is the high yield category) tercile values of the climatological record. For a comprehensive description of the retro-active forecast process, refer to Landman et al. (2012).

We only show Pearson correlation values between cross-validated forecast and observed time series in order to represent the deterministic skill level of the crop model. Two sets of retro-active probabilistic forecasts are used to determine the potential economic value of the crop forecast system (Hagedorn and Smith, 2009). For this purpose we make use of the cumulative profit (CP) values generated by the CPT software. The CP values evaluate probabilistic forecasts by means of quantifying the skill of the forecast using an effective daily interest rate. Some capital is invested into the first of a series of consecutive probabilistic forecasts, say for example ZAR1,000. Depending on the outcome of how well the forecast performs, a return is obtained on the investment. This return is calculated based on 'fair odds' and assuming that the ZAR1,000 is spread across the forecast categories in proportion to the forecast probabilities. This means that for the observed category (above, below or normal), the farmer is returned three times the amount of money invested in that category each year. The CP results can be interpreted as follows: for the CP value of, say, 20 found for a specific retro-active forecast year, it means that an initial investment of ZAR1,000 in the first year would be worth ZAR(1,000x20=)20,000 in the specified year. One would subsequently invest all of the ZAR20,000 on the next year's forecast, and so forth. See Mason (2018) for a comprehensive explanation on the calculation of CP values.

Results

The cross-validation hindcasts and observed values, both normalised here, are presented in Fig 1. On the figure El Niño and La Niña years are respectively shown as "El" and "La". Also presented in the figure are the differences between the observed and predicted crop indices.

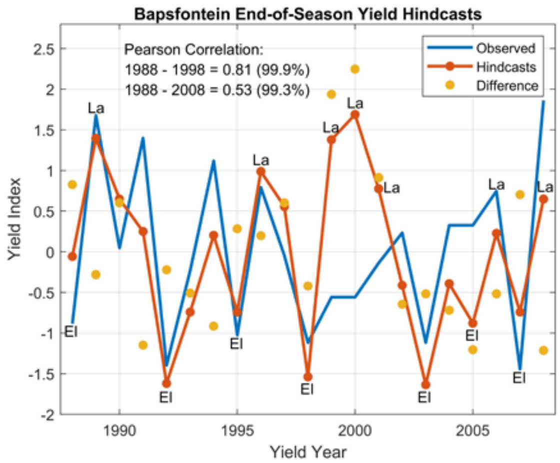

Figure 1. Cross-validation crop-yield hindcasts vs. observed values. The time series have been normalised. The Pearson correlation values for the first 11 years (1988 to 1998) and for all the years are presented on the figure, along with their levels of statistical significance.

Two features of the cross-validated results are immediately apparent. The first is that all the El Niño years are associated with below average predicted yields, and all the La Niña seasons are associated with above average predicted yields. One may thus deduce that ENSO phases play a significant part in this crop model's yield predictions for the Bapsfontein farm. Second, although all but one (2005) of the El Niño years are found to be actual low-yield years, a number of La Niña years are also actual low-yield years. For example, for the three-year period of 1999 to 2001 the observed yields are below average. The reasons for the forecast failures during these three seasons may be related to the observed rainfall outcomes during those summers: the 1998/99 and 2000/01 seasons were not wet La Niña years over parts of SADC leading to low levels of soil moisture, and during 1999/00 the larger part of the region was flooded and possible damages to crops occurred.

Corresponding author: Tel: +27-12-420-3713; Email: Willem.Landman@up.ac.za

The deterministic skill level of the cop-yield model is typical for southern Africa (c.f. Landman et al., 2012). Moreover, the correlation between forecasts and observed for the first half of the test period is very high (correlation > 0.8), suggesting that future forecasts may turn out to be skilful. However, the yield forecasts associated with the three La Niña seasons of 1998/99 to 2000/01 turned out to be poor. The question we want to answer in this study is whether or not these poor yield forecasts could have had a detrimental effect on the finances of a farmer basing decisions on them – specifically investment decisions, for the sake of the example. We address this question by first calculating the CP values for the seasons following the first 11 years of the data, i.e. from 1999 to 2008 as shown in Fig. 2. Negative CP values are found for the majority of the years, with positive values found only for 2005, 2007 and 2008. In fact, it was only at the end of the retro-active period when the farmer obtained substantial and positive CP's. The second retro-active period that excludes the poor forecasts associated with the three consecutive La Niña seasons present much better CP outcomes. In fact, it was primarily only during the 2005 season when there is a big difference between the forecast and observed anomalies where the CP value of this retro-active period is negative. Take note how the inclusion of the poor forecasts in the CP calculations has significantly delayed financial recovery, and how much less of a detrimental effect a single poorly forecast season (2005) has on profits compared to when consecutive seasons of poor forecasts are included. The above is a simple example, which can be expanded to include more realistic investment and planting strategies e.g. including maize prices and input costs.

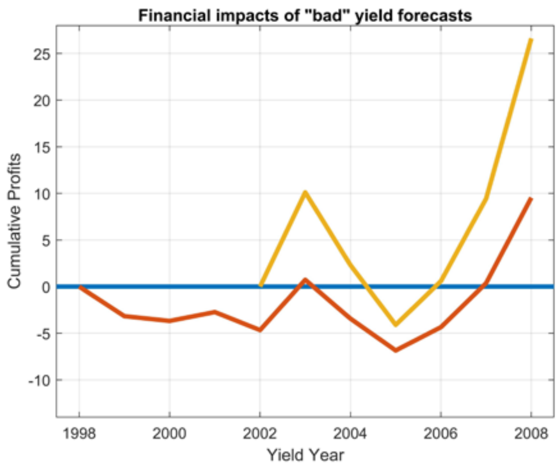

Figure 2. Cumulative profits as determined over the two retro-active forecast periods. Red: results from the retro-active period from 1999 to 2008; Orange: results from the retro-active period from 2003 to 2008.

Summary and Conclusions

Southern Africa ranks poorly against the majority of regions where ENSO has an effect on seasonal-to-interannual climate variability (Landman et al., 2019). Notwithstanding, seasonal forecasts have been found skilful over certain areas of the region and in particular during certain times of the year (Landman et al., 2012; Archer et al., 2019). Moreover, the majority of end-of-season crop yields over areas which include the Bapsfontein farm, are likely to be predictable when there is an ENSO event taking place (Landman and Beraki, 2012). Here we presented a linear statistical crop-yield prediction model that uses output from a coupled climate model that is linked with dry-land crop-yields at a single farm. The results presented may not be representative for all crop farms in the SADC region, since not all end-of-season yields may be equally influenced by ENSO events.

Verification of the hindcasts produced by the crop model shows significant levels of skill. In fact, the first half of the 21-year deterministic verification period shows an unprecedented level of skill (corr > 0.81). Notwithstanding, the model performed poorly for three consecutive seasons during a series of La Niña events. We wanted to find out how such a sequence of poor forecasts may have affected the farmer during these three and subsequent years. The main conclusions that may be derived from this study are that the consecutive poor forecasts could have devastating consequences for the farmer, and that a possible financial recovery may only have happened a good number of years after the three poorly predicted years. We also show that a single poorly predicted season does not necessarily have the same negative financial impact.

So what implications does this result hold for seasonal forecast model developers, as well as those concerned with applying seasonal forecasts? There are several immediate implications, including:

- Even using a skilful seasonal forecast model may not immediately translate into tangible benefits to the user (or farmer as in this case), but may require sustained use of skilful forecasts over a period of several years;
- The yield forecasts and ENSO are symmetrical (low yields predicted during El Niño; high yields predicted during La Niña), but this symmetry is not as evident in the observed outcomes since not all La Niña seasons produced high yields;
- Even the hedging that takes place in CP calculations for each year (the capital is spread across the three categories according to the predicted probabilities) does not

Corresponding author: Tel: +27-12-420-3713; Email: Willem.Landman@up.ac.za

guarantee positive benefits. However, placing all resources on the assumption of single category outcome is even riskier.

It is clear that we may have to shift our priorities towards addressing user needs through tailored forecasting and the honest conveyance of model forecast caveats to users (including communication of uncertainty). Focussing our limited resources on demonstrating our capabilities as a modelling community, to address technical modelling challenges such as the production of high-resolution forecasts, risks ignoring fundamental limitations in using all seasonal forecasts. This is especially concerning in light of recent research which demonstrates that high resolution seasonal forecasts may hold very little benefit (Scaife et al., 2019). There are also some potentially relatively easy gains, such as exploring the benefit of using other prediction variables (e.g. pre-season soil moisture, evaporation-related variables etc), or combining a wider range of publicly available seasonal forecasts.

We need to be much more honest and transparent about our prediction capabilities even with skilful models such as the one presented here. Indeed, communicating how to use and interpret seasonal forecasts, as well as linking them to variables and impacts of interest to specific user groups in particular locations may hold as much, if not more benefit. Importantly, the process of producing responsible seasonal forecasts goes beyond producing the forecast itself and must not undermine trust between forecast producers and users through overblown promises of forecast accuracy and skill. To do so is to risk our efforts being misunderstood and ultimately ignored.

Acknowledgements

The work is a result of collaboration between South Africa and the IRI, supported by the National Research Foundation of South Africa (NRF) through their financial support of rated researchers, and the ACyS-ACCESS project funded by the NRF.

References

Archer, E., Landman W.A., Malherbe, J., Tadross, M. and Pretorius, S. (2019). South Africa's winter rainfall region drought: a region in transition? Climate Risk Management. **25**, 100188. DOI: 10.1016/j.crm.2019.100188.

Beraki, A., DeWitt, D., Landman, W.A. and Olivier, C. (2014) Dynamical seasonal climate prediction using an ocean-atmosphere coupled climate model developed in partnership between South Africa and the IRI. Journal of Climate. **27** 1719–1741. DOI:10.1175/JCLI-D-13-00275.1.

Barnston, A.G. and Tippett, M.K. (2017). Do statistical pattern corrections improve seasonal climate predictions in the North American Multimodel Ensemble models? Journal of Climate. **30**, 8335-8355. DOI: 10.1175/JCLI-D-17-0054.1.

DeWitt, D. G. (2005). Retrospective forecasts of interannual sea surface temperature anomalies from 1982 to present using a directly coupled atmosphere–ocean general circulation model. Monthly Weather Review. **133**, 2972–2995.

Hagedorn, R. and Smith, L.A. (2009). Communicating the value of probabilistic forecasts with weather roulette. Meteorological Applications. **16**, 143–155, DOI: 10.1002/met.92.

Landman, W.A. and Beraki, A. (2012). Multi-model forecast skill for mid-summer rainfall over southern Africa. International Journal of Climatology.**32**, 303–314. DOI: 10.1002/joc.2273.

Landman, W.A. and Goddard, L. (2002). Statistical recalibration of GCM forecasts over southern Africa using model output statistics. Journal of Climate, 15, 2038-2055.

Landman, W.A., Barnston, A.G., Vogel, C. and Savy, J. (2019). Use of El Niño-Southern Oscillation related seasonal precipitation predictability in developing regions for potential societal benefit. International Journal of Climatology, DOI: 10.1002/JOC.6157.

Landman, W.A., DeWitt, D. Lee, D.-E., Beraki, A. and Lötter, D. (2012). Seasonal rainfall prediction skill over South Africa: 1- vs. 2-tiered forecasting systems. *Weather and Forecasting*. **27**, 489-501. DOI: 10.1175/WAF-D-11-00078.1.

Malherbe, J., Landman, W.A., Olivier, C., Sakuma, H and Luo, J.-J. (2014) Seasonal forecasts of the SINTEX-F coupled model applied to maize yield and streamflow estimates over north-eastern South Africa. Meteorological Applications. **21**, 733–742. https://doi.org/10.1002/met.1402.

Mason, S.J. (2018). *Guidance on Verification of Operational Seasonal Climate Forecasts*. World Meteorological Organization, Commission for Climatology XIV Technical Report 1220.

Scaife, A.A., Camp, J., Comer, R., Davis, P., Dunstone, N., Gordon, M., MacLachlan, C., Martin, N., Nie, Y., Renm H.-L., Roberts, M., Robinson, W., Smith, D. and Vidale, P.L. (2019). Does increased atmospheric resolution improve seasonal climate predictions? Atmospheric Science Letters. e922, https://doi.org/10.1002/asl.922.

Wilks, D.S. (2011). Statistical Methods in the Atmospheric Sciences, 3rd edition. Academic Press: Amsterdam, pp. 676.

Corresponding author: Tel: +27-12-420-3713; Email: Willem.Landman@up.ac.za