

Interactive comment on “From skill to value: isolating the influence of end-user behaviour on seasonal forecast assessment” by Matteo Giuliani et al.

Matteo Giuliani et al.

matteo.giuliani@polimi.it

Received and published: 31 March 2020

This is a study into the value of inflow forecasts in water release decision making, focusing on the benefits to agricultural profitability. Previous studies have demonstrated how forecasts can be adopted in reservoir operations to marginally improve on a pre-specified objective. This paper offers an incremental advance by coupling the reservoir model to an agricultural model, allowing for calculation of profits associated with the updated release schedule. The subject is certainly of general interest to the hydrological community. The paper is well written and the method easy to follow. While the study is mostly sound (I have a few concerns outlined below), a significant contribu-

C1

tion to knowledge is missing. One can easily deduce without this study that reduced water supply deficits in a reservoir release model should lead to increased profits for crop-growers relying on that water. The monetary values provided cannot be offered as a contribution either, since they are not reflective of actual profit gains that would be gleaned by crop-growers (partly because the reservoir operations are stylized for this case study rather than representing real world operations). The approach taken is described as a “novel evaluation framework”. It appears to be a forecast product providing information for a reservoir model, the release from which forces an agricultural profit model. One-way coupling of models (which is what I understand this to be) does not qualify as a novel framework in my view. Lastly, because the study is conducted on a single site and using a short time series with only one drought (with much of the analysis drawn from performances during that drought), the conclusions are not generalizable. The authors acknowledge this lack of generalizability in their final sentences, and I think that their suggestions for future research are actually required in this paper to help with the knowledge contribution. While a single case study can be valuable, I cannot see compelling new insights on value of forecasts arising from this analysis to warrant publication. I think the suggested exponential relationship between forecast skill and farmer profit could be a significant contribution if demonstrated and elaborated more carefully through more detailed analysis across a range of possible droughts and with incremental adjustments to the forecast skill. I would be supportive of publication of this paper if the authors can (a) deepen their analysis to generate a more compelling advance on existing knowledge, and (b) address the small number of other concerns listed below.

We agree with the reviewer that our evaluation framework per se may not represent a sufficient contribution to the existing literature. However, in our opinion there are aspects other than the integrated modelling chain that are novel, such as the use of different river flow forecasts as inputs to understand which part of the hydrological modelling chain is relevant in this case, as well as the inclusion of the decision maker’s perspective by looking at specific forecast quantiles. In the revised version of the pa-

C2

per, we will clarify that these two aspects, along with the inference of a relation between gains in forecast skill and gains in end-user profit, represent the main methodological contributions of the paper. Moreover, we respectfully disagree that our quantification of the value of hydroclimatic services in terms of estimates of potential economic benefit to the end-users can be summarized as a “forecast product providing information for a reservoir model, the release from which forces an agricultural profit model”. Our evaluation framework combines a state-of-the-art hydroclimatic service run by SMHI with a detailed model of the Lake Como basin. This latter couples an advanced operational module to simulate the lake operation, including an optimization of the operational decisions via Reinforcement Learning techniques, with an accurate description of the agricultural district provided by a spatially distributed model with a regular mesh of cells with a side length of 250 m. Previous works (e.g., Giuliani et al., 2016) demonstrated that this model is capturing the main characteristics of the real systems, including the actual operations, and its outputs were validated with respect to observed data both in terms of lake releases and of agricultural practices such as water requirements and consumption, crop production, or land use decisions. We therefore consider our estimates to be a valuable contribution for the case study area and in the revised version of the paper, we will better clarify the potential value of our results for the considered case study.

Giuliani, M., Li, Y., Castelletti, A., and Gandolfi, C.: A coupled human-natural systems analysis of irrigated agriculture under changing climate, Water Resources Research, 52, 6928–6947, 2016.

Specific comments: The decision to use quadratic water supply deficit as the objective function is not fully justified. If the purpose of the water supply is to meet farmer needs, and if profit is the goal of the farming community, then why not use farmer profit as the objective? This would greatly improve the interpretability of the results, particularly for aim (iii) “the inference of the relation between gains in forecast skill and gains in end-user profit.” Currently, the paper lacks discussion on how the discontinuity between the

C3

optimization objective and profitability affects the conclusions drawn. In particular, the squaring of the water deficit objective would normally result in hedging behavior that would reduce overall profit to avoid possible cases of extreme deficit. It’s also not clear from this analysis how water deficit affects profit. Does a small deficit necessarily imply loss of crop production, or can farms run at full profitability during periods of small or moderate deficit?

The quadratic water supply deficit is a traditional formulation in reservoir operations since the work by Hashimoto et al. (1982). This objective generates hedging strategies that minimize large deficits, while accepting small, distributed deficits; these strategies are known to be suitable for irrigation practices as crops can resist in case of small deficits while extreme ones can generate crop failures. Obviously, the larger the deficit, the larger will be the difference between potential and actual evapotranspiration computed in the crop-yield module of the agricultural district, with this delta generating large stress and loss of production according to the approach proposed in FAO (Doorenbos et al., 1979) that is implemented in our model. For example, moving along the baseline Pareto front from the policy selected in Figure 4 towards solutions that attain lower deficits (e.g., a policy P1 with squared deficit equal to 2749 (m³/s)² or policy P2 with squared deficit equal to 2672 (m³/s)²) generate higher profits for the farmers, specifically 24.6 M€year for P1 and 24.7 M€year for P2. Moreover, the computation of the water supply deficit includes a time-varying parameter that penalizes more the deficit after germination to the beginning of phenological maturity, with these crop stages determined by the agricultural district model. We will add a sentence to discuss this aspect in the revised version of the paper. The reason for not directly using the farmer profit as an objective function relies in the computational requirements of the agricultural model simulations (which is based on a mesh of about 11,000 cells with a side length of 250 m) that are incompatible with the computational costs of the EMODPS approach used for the design of the optimal Lake Como operations. The EMODPS optimization indeed requires running 40 million simulations for each forecast input, and the overall analysis comprises a total of 320 million simulations that required approximately

C4

42,670 computing hours. We will clarify this point in the revised version of the paper. Lastly, it is worth mentioning that the validation of the model in Giuliani et al. (2016) showed that a policy designed using this formulation generates a good approximation of the observed operations of the lake.

Doorenbos, J., Kassam, A., and Bentvelsen, C.: Yield response to water, irrigation and drainage. Paper no. 33, Tech. Rep., Food and Agriculture Organization, Rome, Italy, 1979.

Hashimoto, T., Stedinger, J., Loucks, D. (1982). Reliability, resilience, and vulnerability criteria for water resource system performance evaluation. Water Resources Research, 18(1), 14–20.

The decision to vary the ensemble selected from mean to 10th and 25th percentiles to capture drought risk aversion requires better justification, too. It would seem more prudent to adjust the objective to represent risk-averse preference (e.g., increasing the exponent applied to the objective, or, if changing the objective function to farmer profit, adding penalties for very significant losses) than to deliberately under-estimate the inflow.

The idea of exploring alternative interpretations of the forecast ensemble by replacing the mean with low percentiles is motivated by the growing literature suggesting that, at seasonal time scales, probabilistic forecasts are often used to convey uncertainties related to initial hydro-climatic conditions, scenarios of predicted meteorological conditions, and adopted models, potentially adding value for decision making (see Georgakakos and Graham, 2008; Cloke and Pappenberger, 2009). At the same time, there is also growing evidence that higher forecast accuracy does not necessarily imply better decisions because of the challenges associated to the human interpretation of forecasts as well as to the communication of probabilistic information (Ramos et al., 2010, 2013; Crochemore et al., 2016). However, at the best of our knowledge, this point has been so far investigated mostly via serious games, interviews, or direct interactions

C5

with decision makers, while our work aims at providing a quantitative analysis of this challenge by simulating how different behavioral attitudes influence the interpretation of the forecast ensemble and ultimately impact on operational decisions and resulting performance. We will better clarify this contribution in the revised version of the paper.

Cloke, H. and Pappenberger, F. (2009), Ensemble flood forecasting: a review, Journal of Hydrology, 375, 613–626

Crochemore, L., Ramos, M., Pappenberger, F., van Andel, S., and Wood, A. (2016), An experiment on risk-based decision-making in water management using monthly probabilistic forecasts, Bulletin of the American Meteorological Society, 97, 541–551

Georgakakos, K. and Graham, N. (2008), Potential benefits of seasonal inflow prediction uncertainty for reservoir release decisions, Journal of Applied Meteorology and Climatology, 47, 1297–1321

Ramos, M., Mathevet, T., Thielen, J., and Pappenberger, F. (2010), Communicating uncertainty in hydrometeorological forecasts: mission impossible?, Meteorological Applications, 17, 223–235

Ramos, M., van Andel, S., and Pappenberger, F. (2013), Do probabilistic forecasts lead to better decisions?, Hydrology and Earth System Sciences, 17, 2219–2232, doi:10.5194/hess-17-2219-2013

Line 81: please clarify what “heavily man-overworked” means (and why its relevant).

We meant that the water resources in the basin are highly regulated by water infrastructures, including 16 Alpine hydropower reservoirs in the upstream part of the catchment that can store about 545 Mm³, which is more than twice the active volume of the lake; Lake Como in the middle, which is a deep glacial lake whose outlet is controlled since 1946 with the twofold primary purpose of water allocation to the downstream users and flood protection along the lake’s shoreline, along with additional interests related to navigation, fishing, tourism, ecosystems; the lake release serves a dense network

C6

of downstream irrigation canals, which convey water to four agricultural districts with a total surface of 1400 km², mostly growing maize. The same releases are also sufficient to feed eight run-of-river hydroelectric power plants. These features are peculiar characteristics of this system, which should not be confused with a natural lake, and the resulting high level of control of the water in the basin is an important factor that motivates the search for more efficient management strategies relying on hydroclimatic services. We will rephrase this sentence in the revised version of the paper.

Line 89: do you mean “most Southerly” point on the lake shoreline, or the “near the outflow” of the lake?

We mean lowest point in terms of elevation, which is the reason why it is the location suffering the most from the floods. Note that the lake outflow are in the other branch of the lake, while the one where Lake Como is located is a dead branch. We will rephrase the original sentence to clarify this point.

Line 257: Why bother with the Pareto analysis if the flood objective effectively becomes a constraint. I don't think the readers of the study need all of the detail of the Pareto analysis if multi-objective optimization is not actually used to generate the key results.

The flood objective is not a constraint in our problem, but we designed the Pareto optimal set of operating policies by using a truly multi-objective approach, namely the Evolutionary Multi-Objective Direct Policy Search (Giuliani et al., 2016b). Since the result is then a set of solutions that explores the tradeoff between flood control and water supply, we used a reference value of flood days only to filter the Pareto optimal solutions and select one policy for each set, attributing their different performance to the different forecasts that they use. However, it is important to notice that the benefit of informing the lake operations with hydroclimatic services is in both the objectives, with the overall Pareto front that shifts toward the bottom-left corner of the objective space. We will better clarify this point in the revised manuscript by reinforcing the tradeoff analysis narrative prior to focus on the selected policies, for example by reporting some

C7

multi-dimensional metrics (e.g., hypervolume indicator) to quantify the improvement of the full Pareto optimal set.

Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., and Reed, P. (2016b), Curses, trade-offs, and scalable management: advancing evolutionary multi-objective direct policy search to improve water reservoir operations, Journal of Water Resources Planning and Management, 142

Line 260: the fact that profits are improved through operations is used to support the idea that forecasts can be valuable for managing extreme drought. Presumably the impact is greatest during drought because this is the only time when profit can be compromised (i.e., average flow conditions are unlikely to lead to supply deficits, meaning forecasts are not actually useful except leading up to and during drought). Is this correct? If so, why not focus analysis on droughts and also introduce other drought events to help support and generalize these conclusions?

In the case of Lake Como, the role of operations is larger than what the reviewer says because the natural water availability (i.e., the lake inflows) is not covering the downstream demand and the system would experience deficits during the summer. This is the reason why hydroclimatic services are expected to be valuable also in normal conditions, and likely also in wet years as they would suggest the operator to keep a larger flood buffer by releasing more water than in normal conditions as high inflow volumes are expected over the upcoming months. We will better clarify the central role of the lake operations in the revised version of the paper.

Line 311: Has this function been fitted across all of the points on Figure 6? Please justify or comment on the appropriateness of combining the all-years and 2005 results in the same function. The idea of exponential relationship between profit and forecast skill would be a powerful conclusion, but is surely best demonstrated using (a) a model that can adjust forecast skill incrementally allowing generation of many data points, and (b) repeating the analysis across multiple droughts.

C8

Yes, the function is fitted across all the points in the Figure. We agree with the reviewer that having more points would make this result more statistically sound. However, as mentioned also in previous replies, the data/modeling/computational requirements of our analysis are not negligible, thus limiting the possibility of easily generating more points. We therefore consider this result anyway acceptable in the context of our paper, where this function is one out of three contributions, and we will clarify the associated limitations of such analysis in the revised version of the manuscript.

Interactive comment on Hydrol. Earth Syst. Sci. Discuss., <https://doi.org/10.5194/hess-2019-659>, 2020.