26th International Congress on Modelling and Simulation, Adelaide, SA, Australia, 30 Nov–4 Dec 2025
mssanz.org.au/modsim2025

Rapid assessment of environmental flow targets using Bayesian logistic regression

Andrew Schepen[a] and **David E. Robertson**[b]

[a] *CSIRO Environment, Dutton Park, Queensland, Australia,* [b] *CSIRO Environment, Clayton, Victoria, Australia*

Abstract: We present a Bayesian logistic regression framework to rapidly estimate the probability of achieving site-specific flow indicators (SFIs) under varying inflows and environmental water recovery levels. Using inflow and environmental water recovery volume as predictors, the model provides reliable probability estimates and quantifies uncertainty. Applied to Barmah–Millewa and Chowilla, it demonstrates strong predictive skill and offers a computationally efficient alternative to full river system modelling for initial exploratory scenario analysis.

Keywords: *Site-specific flow indicators, environmental flows, prediction, uncertainty*

1. INTRODUCTION

Environmental flows are fundamental to sustaining riverine and floodplain ecosystems, but their availability is constrained by high climate variability and competition from consumptive water use. Within the Murray–Darling Basin, the Basin Plan defines a set of site-specific flow indicators (SFIs; MDBA (2011)) that quantify ecological water requirements at key locations. Each SFI specifies a flow event, characterised by the duration flow exceeds a minimum flow magnitude, that represents conditions necessary to achieve outcomes such as floodplain forest inundation, wetland connectivity, and habitat provision for native fauna. For example, in the Barmah-Millewa Forest, SFIs include events of 12.5 GL/day for 70 days and 60GL/day for 14 days, clearly assessing different regimes of environmental flows. The Basin Plan specifies target frequencies for achieving SFIs that are based on long modelling results. Whether these SFIs are achieved in any given year depends on the interplay between climatic inflows, storage dynamics, and water recovery for the environment after diversions for agriculture and other consumptive uses. As a result, assessing the likelihood of meeting SFIs in any given year requires tools that integrate hydrological variability, water management rules, and ecological objectives.

One approach to assessing the effect of environmental water recovery, climate variability and change on SFI achievement is through full river system modelling. However, such modelling is computationally intensive and usually undertaken for stationary climate and environmental water recovery conditions, e.g. no recovery of environmental water or 2750 GL recovery. In reality, environmental water recovery since the adoption of the Basin Plan has been dynamic, with water recovery progressively increasing over time. In addition, inherent climate variability means that frequency of SFI achievement over relatively short periods, such as decades, may differ considerably from the target long-term average. Complementary methods that can rapidly estimate the probability of meeting SFI achievement under dynamic water recovery and climate conditions, may therefore assist in understanding the impact of the Basin Plan since its implementation and also support expedited policy exploration under different climate and water recovery scenarios.

2. METHODOLOGY

We develop a Bayesian logistic regression framework to estimate the annual probability of achieving SFIs. The predictors are system-scale annual inflow and the volume of environmental water recovery, while the binary outcome of meeting each SFI forms the response. Model parameters are estimated in PyMC, which provides full posterior distributions, thereby explicitly quantifying uncertainty in SFI attainment probabilities. The training data is restricted to the two benchmark recovery scenarios (0 or 2750 GL/year). However, in predictive mode the amount of recovered environmental water is allowed to take any intermediate value. Model fitting and validation use the long-term historical inflow record (1896–2009), with a 5-fold cross-validation scheme to assess predictive skill. To favour parsimony, we experiment with shared-parameter structures in which certain coefficients (e.g., inflow effects) are shared across indicators within a site, while SFI-specific intercepts capture differences in ecological thresholds (see Table 1). The model is applied to two locations, the Barmah-Millewa Forest, and Chowilla.

3. RESULTS

High predictive skill is demonstrated at both locations, with Brier skill scores exceeding 0.6 in cross-validation (e.g., Figure 1, for Chowilla) and the 3Pred-Shared model attractive as a compromise between performance and parsimony. For Chowilla, the selected 3Pred-Shared model demonstrates strong reliability and discrimination. Reliability diagrams show close agreement between predicted probabilities and observed frequencies across most SFIs. ROC analysis indicates excellent discrimination across both sites, with AUC values typically in the range 0.90–0.99, confirming the model's ability to distinguish between achievement and non-achievement. Overall, the results demonstrate that the Bayesian framework provides reliable probability estimates and robust discriminative skill across contrasting ecological sites, supporting its application under variable environmental water recovery conditions. The models are further tested on independent data for the 2013–2024 period, enabling the evaluation of the impacts of environmental water recovery on SFI achievement. Differences in predicted SFI achievement between actual, baseline, and full recovery scenarios demonstrate the model's ability to generalise and highlight the sensitivity of ecological outcomes to the amount of environmental water recovered and climate variability.

Table 1: Name and configuration of tested models in terms of predictors and shared parameters. 'Índependent' means the parameters are independent for each indicator, and 'Shared' means the parameter is the same for all indicators

Name	Predictors	Intercept	Slopes
2Pred-Ind	Annual inflow, ewater	Independent	Independent
2Pred-Shared	Annual inflow, ewater	Independent	Shared
3Pred-Shared	Annual inflow, ewater, lagged annual inflow	Independent	Shared

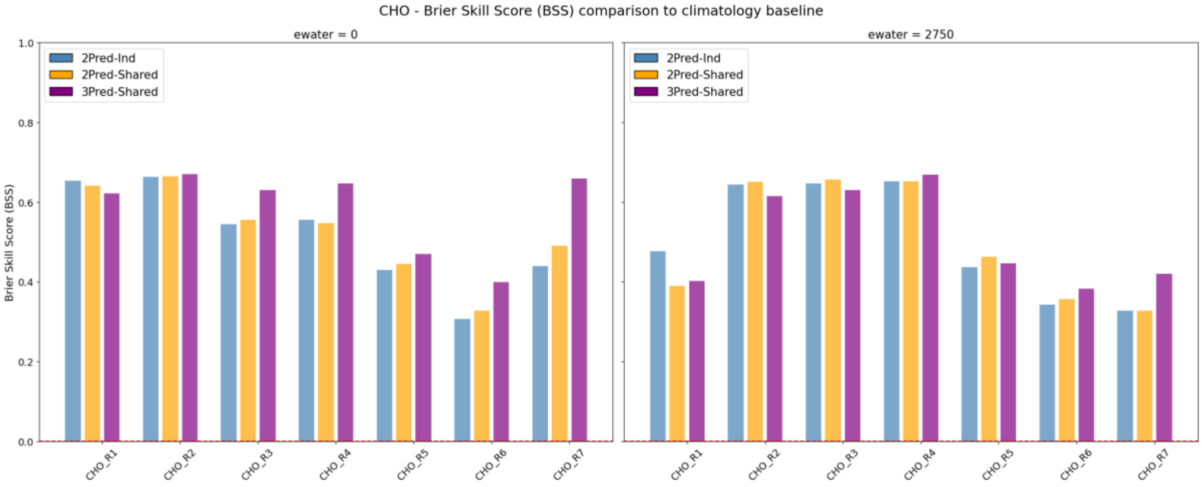

Figure 1: Brier skill scores for the skill of predicting SFI achievement at Chowilla, conditioned on each site-specific flow indicator and fixed environmental water volume

4. CONCLUSION

We developed a Bayesian logistic regression framework for predicting SFI achievement at Barmah–Millewa Forest and Chowilla. The recommended model balances parsimony and predictive skill by using inflow, lagged inflow, and environmental water recovery as predictors, with shared slope parameters across indicators. The model provides reliable probability estimates and quantifies predictive uncertainty. Out-of-sample tests (2013–2024) confirm the model's ability to predict both realised and counterfactual SFI outcomes under variable recovery conditions. This approach offers a computationally efficient alternative to full river system modelling and has potential to support adaptive policy and planning for environmental flows in the Murray–Darling Basin.

REFERENCES

MDBA (2011). "The Proposed 'Environmentally Sustainable Level of Take' for Surface Water of the Murray-Darling Basin: Methods and Outcomes." Murray-Darling Basin Authority, Canberra, Australia.