Water Resources Research

RESEARCH ARTICLE

10.1002/2013WR014185

Key Points:

- Postprocessor approach can be more robust in practice than joint inference
- Joint inference of water balance and autocorrelation parameters is problematic
- Inferring autocorrelation affects parameters and improves aspects of predictions

Correspondence to:

D. Kavetski,
dmitri.kavetski@adelaide.edu.au

Citation:

Evin, G., M. Thyer, D. Kavetski, D. McInerney, and G. Kuczera (2014), Comparison of joint versus postprocessor approaches for hydrological uncertainty estimation accounting for error autocorrelation and heteroscedasticity, *Water Resour. Res.*, 50, 2350–2375, doi:10.1002/2013WR014185.

Received 26 MAY 2013

Accepted 6 FEB 2014

Accepted article online 12 FEB 2014

Published online 18 MAR 2014

Comparison of joint versus postprocessor approaches for hydrological uncertainty estimation accounting for error autocorrelation and heteroscedasticity

Guillaume Evin¹, Mark Thyer¹, Dmitri Kavetski¹, David McInerney¹, and George Kuczera²
¹School of Civil, Environmental and Mining Engineering, University of Adelaide, Adelaide, South Australia, Australia,

²School of Engineering, University of Newcastle, Callaghan, New South Wales, Australia

Abstract The paper appraises two approaches for the treatment of heteroscedasticity and autocorrelation in residual errors of hydrological models. Both approaches use weighted least squares (WLS), with heteroscedasticity modeled as a linear function of predicted flows and autocorrelation represented using an AR(1) process. In the first approach, heteroscedasticity and autocorrelation parameters are inferred jointly with hydrological model parameters. The second approach is a two-stage “postprocessor” scheme, where Stage 1 infers the hydrological parameters ignoring autocorrelation and Stage 2 conditionally infers the heteroscedasticity and autocorrelation parameters. These approaches are compared to a WLS scheme that ignores autocorrelation. Empirical analysis is carried out using daily data from 12 US catchments from the MOPEX set using two conceptual rainfall-runoff models, GR4J, and HBV. Under synthetic conditions, the postprocessor and joint approaches provide similar predictive performance, though the postprocessor approach tends to underestimate parameter uncertainty. However, the MOPEX results indicate that the joint approach can be nonrobust. In particular, when applied to GR4J, it often produces poor predictions due to strong multiway interactions between a hydrological water balance parameter and the error model parameters. The postprocessor approach is more robust precisely because it ignores these interactions. Practical benefits of accounting for error autocorrelation are demonstrated by analyzing streamflow predictions aggregated to a monthly scale (where ignoring daily-scale error autocorrelation leads to significantly underestimated predictive uncertainty), and by analyzing one-day-ahead predictions (where accounting for the error autocorrelation produces clearly higher precision and better tracking of observed data). Including autocorrelation into the residual error model also significantly affects calibrated parameter values and uncertainty estimates. The paper concludes with a summary of outstanding challenges in residual error modeling, particularly in ephemeral catchments.

1. Introduction

1.1. General Motivation for Improving Residual Error Models

Improving calibration and predictive use of rainfall-runoff models is an important objective in hydrological sciences and engineering. Research on probabilistic uncertainty quantification has proceeded along several general directions, including: (1) improving aggregated approaches that aim to characterize the total uncertainty using residual error models (e.g., Sorooshian and Dracup [1980], Kuczera [1983], Schoups and Vrugt [2010], Evin *et al.* [2013], and others) (2) developing more detailed disaggregation approaches that attempt to characterize individual sources of error (e.g., Kavetski *et al.* [2003], Vrugt *et al.* [2005], Reichert and Mieleitner [2009], Renard *et al.* [2011], and others).

Our focus in this study is on aggregated approaches, which, in general, are conceptually simpler and less data intensive than disaggregation approaches. They are hence of particular relevance to operational applications [e.g., Tuteja *et al.*, 2011] and studies over large numbers of catchments [e.g., Perrin *et al.*, 2001]. Aggregated approaches generally rely on a residual error model to provide a statistical description of the differences between the model predictions and observed data (without trying to disentangle its contributing sources). Adequate representations of residual errors are therefore crucial for obtaining reliable and precise hydrological predictions. If poor assumptions are made about the properties of residual errors poor parameter estimates and unreliable predictions are obtained [e.g., Beven and Binley, 1992; Thyer *et al.*, 2009]. Uncertainty analysis, both of parameter estimates and streamflow predictions, is particularly important

when hydrological models are applied outside the range of calibration data, e.g., during forecasting, regionalization, etc. [Zhang *et al.*, 2008].

We also note the broader debates on the merits of “formal” versus “informal” uncertainty estimation methods [e.g., Mantovan and Todini, 2006; Beven *et al.*, 2008]. From our perspective, the ad hoc nature of “informal” methods such as Generalized Likelihood Uncertainty Estimation (GLUE) makes it difficult to systematically develop, test, and improve descriptions of predictive uncertainty. We refer the interested reader to Stedinger *et al.* [2008] and the recent exchange between Clark *et al.* [2012] and Beven *et al.* [2012] for detailed discussions.

1.2. Treatment of Heteroscedasticity: Transformations Versus Direct Representations

It has long been known that residual errors in hydrological applications are generally heteroscedastic and autocorrelated [Sorooshian and Dracup, 1980]. A variety of strategies have been proposed to deal with these features. For example, heteroscedasticity is often taken into account by transforming observed and predicted streamflows using logarithmic or Box-Cox transformations before applying constant-variance statistical error models [e.g., Kuczera, 1983; Bates and Campbell, 2001; Morawietz *et al.*, 2011]. However, Box-Cox and similar transformations can be seen to represent an indirect treatment of heteroscedasticity—their primary effect is to transform non-Gaussian distributions, especially skewed ones, to near-Gaussian shape. Therefore, this paper focuses on more direct representations of error heteroscedasticity, in particular, on strategies where the variance of the residual errors is conditioned directly on the runoff magnitude. The direct conditioning strategy has been used in a number of studies [e.g., Sorooshian and Dracup, 1980; Thyer *et al.*, 2009; Schoups and Vrugt, 2010; Renard *et al.*, 2011]. Separate research beyond this paper is needed to more clearly establish the advantages and disadvantages of direct versus transformational approaches, in particular, for capturing the skew and kurtosis of residual errors, avoiding negative flows, ensuring stable predictions, etc.

1.3. Treatment of Error Autocorrelation

Streamflow residual errors are generally autocorrelated, most notably due to the propagation of input and structural errors through model storage components, which induce “memory” effects in the error time series [Kavetski *et al.*, 2003]. The autocorrelation of residual errors can be represented using a general autoregressive moving average (ARMA) model [Kuczera, 1983], though Schaeffli *et al.* [2007] report that a simpler autoregressive (AR) model of order 1 appears sufficient.

1.4. Estimation of Residual Error Parameters: Joint Versus Postprocessor Approaches

The parameters of residual error models in hydrological applications are seldom known a priori and therefore must be inferred as part of the calibration. The challenges of jointly estimating parameters that represent persistence and random noise are well known in the stochastic hydrological community; various techniques including conditional approaches that combine the method of moments and maximum-likelihood estimation have been investigated [Stedinger *et al.*, 1985]. These challenges have received less attention in rainfall-runoff modeling, thus forming one of the motivations for this work.

In this study, we distinguish between two inference approaches. The first is a “joint inference” approach, where the residual error model parameters are estimated simultaneously with the hydrological model parameters. The second is a “postprocessor” approach, where inference is undertaken in two or more consecutive stages. For example, in Stage 1, the hydrological model parameters are calibrated using a simple residual error model, followed by Stage 2 where the parameters of a more complex residual error model are estimated keeping the hydrological parameters fixed at the values estimated in Stage 1.

The joint approach corresponds to standard Bayesian and maximum-likelihood approaches and, in the context of inferring the error autocorrelation, has been used by Kuczera [1983], Bates and Campbell [2001], Schoups and Vrugt [2010], and others. Theoretically, joint inference provides a full assessment of predictive and parameter uncertainty, including the effects of all parameter interactions. However, strong interactions can lead to poor parameter identifiability and to numerical difficulties such as slow convergence of optimization and sampling algorithms [Renard *et al.*, 2010; Evin *et al.*, 2013].

The postprocessor approach is popular in forecasting applications [e.g., Engeland *et al.*, 2010; Morawietz *et al.*, 2011; van Andel *et al.*, 2013], and has been used in the recent scheme by Pianosi and Raso [2012]. Montanari and Koutsoyiannis [2012] also advocate an approach where the predictive error model is estimated

conditionally on the hydrological model parameters. This approach can be viewed as an “uncertainty postprocessor” and is computationally appealing due to the lower number of parameters that are estimated at any stage of the calibration. However, unlike the joint approach, the postprocessor strategy ignores interactions between hydrological parameters and residual error parameters, and produces estimates of residual error parameters that are conditional on a particular point estimate of hydrological parameters. Moreover, postprocessor methods use an “incomplete” error model when calibrating the hydrological parameters (e.g., Stage 1 of the calibration may omit the error autocorrelation). Due to these omissions, parameter estimates obtained using a postprocessor approach can be unreliable (biased and/or with poor uncertainty estimates), especially if overly simplistic error models are used in Stage 1. This may have implications for parameter and model interpretability.

It can be seen that joint and postprocessor approaches have a number of distinct advantages and limitations, from both theoretical and practical perspectives. However, due to a lack of systematic comparisons between these inference strategies, it is currently difficult to appraise their relative strengths and weaknesses in terms of predictive performance and parameter estimation, and make practical recommendations for general scientific and operational applications in hydrology.

1.5. Study Aims and Objectives

This study evaluates and compares calibration strategies and error models for describing the autocorrelation and heteroscedasticity of residual errors of hydrological models. The specific aims are:

1. Compare joint versus postprocessor inference strategies for describing the autocorrelation of residual errors when calibrating hydrological models. The comparison focuses on parameter estimation and streamflow prediction, and is pursued using both synthetic data and real data from a range of catchments, and two hydrological models. The following questions are explored:

- a. Does the performance of the inference strategy vary with the hydrological model?
- b. Does the performance of the inference strategy vary with the catchment?

Comparison across different hydrological models is important because different types of interactions can occur between hydrological and error model parameters depending on the model structure and parameterization. For example, *Evin et al.* [2013] illustrate how heteroscedasticity and autocorrelation parameters can interact strongly with a hydrological model parameter controlling the water balance.

Model performance also varies substantially from one catchment to another. For example, calibration and predictive uncertainty quantification are much more challenging in ephemeral catchments [Ye et al., 2008; Smith et al., 2010]. Hence, as detailed in section 3.2, we use catchments with a wide range of hydroclimato-logical regimes, from very wet to dry (ephemeral).

2. Evaluate the benefits of accounting for the autocorrelation of residual errors when making hydrological predictions. We compare the predictions obtained using residual error models with and without autocorrelation. Similar to Aim 1, this comparison investigates impacts on parameter uncertainty estimates and streamflow predictive distributions. The evaluation also considers the impact on time aggregated predictions, e.g., when aggregating daily streamflows to the monthly scale, as is necessary in seasonal streamflow forecasting applications [e.g., Tuteja et al., 2011].

The paper is structured as follows. Section 2 describes the residual error models used in the empirical comparisons. Section 3 describes materials (data, hydrological models) and diagnostic tools (metrics, classification). Sections 4 and 5 present the results obtained with synthetic and real data, respectively, followed by a discussion and synthesis in section 6. The paper concludes with section 7, where practical recommendations and suggested areas for future work are outlined.

2. Description of Residual Error Models

2.1. Bayesian Methods for Residual Error Modeling

A hydrological model predicts streamflow Q as a function of forcing data X (e.g., rainfall and evapotranspiration), hydrological model parameters θ_H , and initial conditions S_0 .

Consider a residual error model based on additive errors ϵ , defined at time step t as

$$\varepsilon_t = \tilde{Q}_t - Q_t^{\theta_H}(\tilde{\mathbf{X}}_{1:t}, \mathbf{S}_0) \quad (1)$$

where $\tilde{Q}_t$ is the observed flow, and $Q_t^{\theta_H}(\tilde{\mathbf{X}}_{1:t}, \mathbf{S}_0)$ is the flow simulated by the hydrological model with parameters θ_H and forced with the observed inputs over time steps 1 to t , $\tilde{\mathbf{X}}_{1:t}$.

Bayes equation yields the posterior distribution of estimated hydrological parameters θ_H and residual error model parameters θ_ε as follows,

$$p(\theta_H, \theta_\varepsilon | \tilde{\mathbf{X}}, \tilde{\mathbf{Q}}) \propto p(\tilde{\mathbf{Q}} | \theta_H, \theta_\varepsilon, \tilde{\mathbf{X}}) p(\theta_H, \theta_\varepsilon) \quad (2)$$

The likelihood function in equation (2) is given by the joint pdf of the residuals

$$p(\tilde{\mathbf{Q}} | \theta_H, \theta_\varepsilon, \tilde{\mathbf{X}}) = p(\varepsilon | \theta_H, \tilde{\mathbf{X}}, \tilde{\mathbf{Q}} | \theta_\varepsilon) \quad (3)$$

where $\varepsilon | \theta_H, \tilde{\mathbf{X}}, \tilde{\mathbf{Q}}$ is the vector of residual errors, computed over a calibration period comprising N_t time steps. A warm-up period is used to avoid the impact of unknown initial conditions.

2.2. Representation of Heteroscedasticity in the Residual Error Model

Let us define a normalized residual error η_t as follows:

$$\eta_t = \frac{\varepsilon_t}{\sigma_t} \quad (4)$$

where σ_t is a normalization term related to the standard deviation of the raw residual errors, $\sigma_{\varepsilon(t)}$.

In this study, we consider a linear heteroscedastic model conditioned on the simulated streamflow,

$$\sigma_t = \alpha_0 + \alpha_1 Q_t^{\theta_H} \quad (5)$$

The residual error model in equations (4) and (5) introduces the error parameters $\theta_\varepsilon = \{\alpha_0, \alpha_1\}$. In most cases, α_0 is small and serves largely as a safeguard to prevent a collapse of uncertainty (and related difficulties in evaluating the likelihood function) for near-zero flows. Parameter α_1 describes the heteroscedasticity of the errors and typically makes a much larger contribution to predictive uncertainty.

2.3. Calibration Schemes

This study considers the following calibration schemes within the general framework of section 2.1:

1. *WLS*. This scheme is used as a “reference” method. It applies the linear heteroscedastic error model in equation (5) and assumes that the normalized residuals η_t are Gaussian, independent in time and have a unit variance, $\eta \sim N(0, 1^2)$. Under these assumptions, the raw residuals ε_t are independent heteroscedastic Gaussian deviates with time varying standard deviation $\sigma_{\varepsilon(t)} = \sigma_t$. The residual error model parameters $\theta_\varepsilon = \{\alpha_0, \alpha_1\}$ are inferred jointly with the hydrological model parameters.

The streamflow predictive distribution is obtained as follows: (i) sample a set of hydrological and error model parameters, (ii) run the hydrological model, (iii) add a realization from the residual error model to the streamflow time series from step (ii), and repeat steps (i)–(iii) over multiple parameter samples.

2. *WLS-AR1*. This calibration scheme uses the same linear treatment of heteroscedasticity as the WLS approach above, but relaxes the assumption that the errors are independent. Instead, the normalized errors η_t are assumed to be described by a homogeneous AR(1) model with standard Gaussian innovations y_t and an autoregressive parameter ϕ

$$\eta_t = \phi \eta_{t-1} + y_t; \quad y_t \sim N(0, 1^2) \quad (6)$$

As shown by *Evin et al.* [2013], normalizing the residuals before applying the AR(1) model leads to a more robust calibration scheme.

Under the assumption that the innovations y_t have a unit variance, we obtain the likelihood function

$$p(\tilde{\mathbf{Q}}|\theta_H, \theta_e, \tilde{\mathbf{X}}) = p(\varepsilon|\theta_H, \tilde{\mathbf{X}}, \tilde{\mathbf{Q}}|\alpha_0, \alpha_1, \phi) = \left[\prod_{t=1}^{N_t} \sigma_t^{-1} \right] N\left(\eta_1|0, \frac{1}{1-\phi^2}\right) \prod_{t=2}^{N_t} N(y_t|0, 1^2) \quad (7)$$

where $N(x|\mu, \sigma^2)$ is the pdf of a Gaussian random deviate x with mean μ and standard deviation σ .

This calibration scheme is a simple reparameterization of "Approach 2" in *Evin et al.* [2013], with all variances rescaled by a factor $1/(1-\phi^2)$. This reparameterization reduces the strong interactions between parameters ϕ and α_1 noted by *Evin et al.* [2013] and facilitates parameter estimation.

From equation (6) the standard deviation of normalized residual errors is $\sigma_\eta = 1/\sqrt{1-\phi^2}$, and from equation (4) the standard deviation of the raw residual at time step t is

$$\sigma_{\varepsilon(t)} = \frac{\sigma_t}{\sqrt{1-\phi^2}} = \frac{\sigma_0}{\sqrt{1-\phi^2}} + \frac{\sigma_1}{\sqrt{1-\phi^2}} Q_t^{\theta_H} = \alpha_0^* + \alpha_1^* Q_t^{\theta_H} \quad (8)$$

The rescaling to α^* is used when reporting the results for calibration schemes that infer the error autocorrelation, because it enables a direct comparison of α values inferred using all schemes.

The parameters of the WLS-AR1 error model, $\theta_e = \{\alpha_0, \alpha_1, \phi\}$, are inferred jointly with the hydrological parameters (i.e., this scheme represents a fully joint inference). The streamflow predictive distribution is obtained using the same sequence of steps as listed above for WLS.

3. *WLS-AR1-PP*. This is a two-stage calibration scheme: Stage 1 applies the WLS scheme, followed by Stage 2 which estimates the parameters of the WLS-AR1 error model while keeping the hydrological parameters fixed at the optimal ("modal") WLS values. The PP suffix denotes "postprocessor," which is how similar approaches are often described in the hydrological forecasting literature (section 1.4).

The predictive distribution is obtained as follows: (i) run the hydrological model using the optimal hydrological parameters $\theta_H^{(WLS)}$ inferred in Stage 1, (ii) sample a set of error model parameters $\theta_e = \{\alpha_0, \alpha_1, \phi\}$ from the posterior obtained in Stage 2, and (iii) add a realization from the residual error model to the streamflow time series produced in step (i), and repeat steps (ii) and (iii).

The WLS-AR1-PP predictive distribution is a partial predictive distribution because it is conditional on a particular fixed set of hydrological parameters. In contrast, the total predictive distributions for WLS and WLS-AR1 propagate the full posterior uncertainty in the hydrological parameters. Unless otherwise noted, the predictive distributions in this study are marginal in the sense of not making use of newly available observations in the prediction period (as possible in short to mid term forecasting).

Note that residual error models based on equation (1) do not distinguish between observational versus structural components of the total predictive uncertainty—they simply describe differences between observations and predictions. Therefore, unless a separate decomposition of uncertainty is undertaken [e.g., *Renard et al.*, 2011], predictive distributions obtained using such error models represent predictions of streamflow observations $\tilde{Q}_t^{(pred)}$ rather than of actual ("true") streamflows Q_t . Despite the fundamental interest in the prediction of actual quantities, the modeling methodology is consistent inasmuch as we are comparing "predicted" observations with the actual measured data. Indeed, the inability to obtain exact measurements generally limits practical verifications of environmental predictions to comparisons with observed data. This is the case for the empirical comparisons undertaken in this work, where no uncertainty decomposition is attempted. Henceforth, to avoid unwieldy notation, we use $Q_t \equiv \tilde{Q}_t^{(pred)}$ when referring to the streamflow predictions reported in this study. Further technical details of the methodology used to generate the streamflow predictions are provided in Appendix A.

2.4. Handling Potential Negative Flows in Prediction

Unless the assumed distribution of residuals is bounded or truncated, predicted streamflows can be negative, which is clearly unrealistic. To address this problem, in this study, the predictive streamflow distribution is truncated at zero. Algorithmically, this is implemented by sampling the residual errors from the assumed residual error model with inferred parameters, but resampling negative streamflows when they occur. We also considered an alternative mixed "continuous-discrete" truncation approach, where negative streamflows are set to zero. However, this produces a "spike" at $Q = 0$ and can degrade predictive reliability. Although the technical aspects of the zero-truncation procedure require further improvement, the simple

Table 1. Summary of the Hydroclimatology of the MOPEX Catchments^a

No	Name	USGS ID	State	Area (km ²)	Elevation (m)	Mean Rainfall (mm)	Mean Runoff (mm)	Runoff Coefficient (-)	Mean Snow (mm)	Mean Winter Temperature (°C)	Fraction of Daily Runoff Below 1 mm (-)
B1	South Branch Potomac	01608500	WV	3810	171	1055	341	0.32 [8]	2582 [1]	0.3	0.74 [5]
B2	Monocacy River	01643000	MD	2116	71	1050	420	0.40 [4]	1737 [4]	1.3	0.69 [8]
B3	Rappahannock River	01668000	VA	4134	17	1037	377	0.36 [7]	733 [7]	3.0	0.70 [7]
B4	Tygart Valley River	03054000	WV	2372	390	1198	735	0.61 [1]	1901 [2]	-0.3	0.47 [11]
B5	Bluestone River	03179000	WV	1020	465	1036	416	0.40 [3]	798 [6]	2.0	0.67 [9]
B6	East Fork White River	03364000	IN	4421	184	1013	377	0.37 [6]	679 [8]	-0.5	0.71 [6]
B7	French Broad (Wettest)	03451500	NC	2448	594	1413	800	0.56 [2]	902 [5]	3.6	0.16 [12]
B8	English River	05455500	IA	1484	193	900	269	0.29 [9]	1776 [3]	-2.9	0.83 [3]
B9	Spring River	07186000	MO	3015	254	1075	298	0.27 [10]	377 [9]	3.0	0.83 [4]
B10	Amite River (Wet)	07378500	LA	3315	0	1563	609	0.38 [5]	25 [10]	11.2	0.66 [10]
B11	Guadalupe (Driest)	08167500	TX	3406	289	763	116	0.15 [12]	7 [11]	10.1	0.95 [1]
B12	San Marcos River (Dry)	08172000	TX	2170	98	825	179	0.21 [11]	4 [12]	11.2	0.92 [2]

^aThe numbers in the square brackets indicate the ranking according to the respective criterion. The degree of ephemerality is given by the fraction of runoff below 1 mm (last column).

“resampling” approach is deemed sufficient for the purposes of this study, because it fulfils the key aim of avoiding negative flows (and hence improves the reliability of the predictions), and is applied in an identical way to all calibration schemes under consideration. We also note that the fraction of total replicates (including at the same time step) where flows had to be resampled is generally low, except for the case of clearly poor/failed calibrations (see the performance classification in section 5.1).

3. Case Study Setup and Methodology

3.1. Synthetic Data

As part of Aim 1, the theoretical properties of the calibration schemes are examined using synthetic data. The synthetic data are generated using the GR4J model (section 3.3) as follows:

1. *Set “exact” parameters.* In order to be broadly “realistic,” synthetic parameter values are obtained by calibrating the GR4J model to the observed French Broad River catchment data (see section 3.2) using the calibration scheme WLS-AR1. This procedure provides the “exact” values of hydrological parameters θ_H and error parameters $\tilde{\theta}_e$ to be used in the synthetic experiment.
2. *Generate “exact” streamflow.* A long time series (38 years) of synthetic daily runoff is generated using GR4J forced with the observed French Broad River rainfall data and with parameters $\tilde{\theta}_H$.
3. *Generate “observation” errors.* A 38 year long realization from a heteroscedastic AR(1) error model (WLS-AR1) is generated using error model parameters $\tilde{\theta}_e$. Residual errors leading to negative synthetic “observed” streamflows (Step 4) are resampled when they occur.
4. *Generate “observed” streamflow.* Synthetic “exact” streamflow (Step 2) is corrupted with synthetic “observation” errors (Step 3) to yield the synthetic “observed” data.
5. *Generate multiple replicates.* Steps 3 and 4 are repeated 100 times, thus generating 100 synthetic replicates of “observed” streamflow series. These time series are then split into calibration and validation periods.

Although the notion of “exact” environmental quantities has been rightfully questioned in a number of discussions [Montanari, 2007], the insights from synthetic tests are essential to establish the fundamental properties of the methods and to help interpret their practical performance.

We note that these synthetic data are representative of a wet catchment (see description of the French Broad River below). Qualitatively similar results were obtained when synthetic data were generated using parameters and forcing data from a dry catchment (results not shown).

3.2. MOPEX Data

The practical properties of the calibration schemes are examined using the set of 12 catchments from the Model Intercomparison Experiment (MOPEX) data set. The MOPEX catchments are a set of catchments in the continental USA and include locations with diverse hydroclimatological and physical attributes (see

Duan *et al.* [2006] for details). Table 1 summarizes the hydroclimatology of the MOPEX catchments, with additional description provided next.

Catchment B7 (French Broad River) is particularly wet; it has the highest annual runoff (800 mm), a high runoff coefficient (0.56), and a very small proportion of low flows (only 16% of runoff values below 1 mm). Catchment B4 (Tygart Valley River) has similar properties, though to a lesser extent.

Catchments B1 (South Branch Potomac), B2 (Monocacy River), B4, and B8 (English River) are quite “snowy”—they have, on average, more than 1 m of snow per year. Catchments B1–B9 have relatively cold winters (average temperature from December to March below +4°C) and regular snow events (more than 300 mm per year on average).

Catchments B10 (Amite River), B11 (Guadalupe River), and B12 (San Marcos River) have mild winters and very few snow events (less than 30 mm per year on average).

Catchments B11 and B12 are the driest from the MOPEX set (annual streamflow below 200 mm) and can be classified as “ephemeral” (more than 90% of daily streamflow below 1 mm). Catchments B8 and B9 have similar properties, though to a lesser extent (annual streamflow below 300 mm, more than 80% of daily streamflow below 1 mm).

For all catchments, two daily data records are used: (1) an 8 year calibration period (9 September 1973 to 26 November 1981), preceded by an 8 year warm-up period; (2) a 17 year validation period (27 November 1981 to 1 May 1998), with the preceding 8 year calibration period used as a warm-up. In snow-affected catchments, the rainfall data are augmented with snowmelt estimates obtained using the US National Weather Service SNOW-17 model [Clark *et al.*, 2008]. Adjusted daily estimates of potential evapotranspiration are generated using the Sacramento model [Clark *et al.*, 2008].

3.3. Hydrological Models

Two lumped daily conceptual rainfall-runoff models, GR4J and HBV, are used.

GR4J is a general rainfall-runoff model that was empirically developed to provide, on average, a good “compromise” performance across catchments with a wide range of climatic and hydrological regimes [Perrin *et al.*, 2003]. GR4J has a simple structure representing interception, infiltration, percolation [see Figure 1 in Perrin *et al.*, 2003] and is usually calibrated using four parameters (Table 2). An important feature of GR4J is its ability to allow for groundwater exchanges (both imports and exports) using its parameter θ_2 . The behavior of this parameter, which modifies the water balance of the hydrological model, is of particular interest in this study.

In addition to GR4J, a modified HBV model was used. It is based on the original HBV model [Bergström, 1995; Seibert, 1997], and has three conceptual reservoirs and eight calibrated parameters, but excludes the routing store (Table 3). Computationally, it is implemented using robust time stepping schemes [Clark and Kavetski, 2010]. An important difference between the GR4J and HBV models is the absence of a groundwater exchange flux in HBV.

3.4. Calibration Methodology

The prior distribution $p(\theta_H, \theta_e)$ in equation (2) is defined by uniform priors on all inferred quantities, with parameter ranges specified in Tables 2 and 3.

The most probable (modal) posterior parameters are obtained by maximizing the joint posterior distribution in equation (2). To achieve this, we apply a quasi-Newton algorithm with 1000 multistarts uniformly seeded across the parameter domain [Kavetski and Clark, 2010]. The high number of multistarts reduces the possibility of the optimization becoming trapped at local optima.

Table 2. Parameters of the GR4J Model

Symbol	Name	Units	Lower Bound	Upper Bound
θ_1	Maximum capacity of the production store	mm	100	20,000
θ_2	Groundwater exchange coefficient	mm/day	−100	100
θ_3	Maximum capacity of the routing store	mm	1	500
θ_4	Time base of unit hydrograph UH1	days	0.5	10

Table 3. Parameters of the HBV Model

Symbol	Name	Units	Lower Bound	Upper Bound
S_{FC}	Maximum soil moisture content	mm	1	10,000
γ_{PWP}	Parameter of ET process	mm	1	10,000
β	Shape parameter in soil store	-	1	50
L	Threshold depth in fast store	mm	0.5	50
k_0	Parameter of upper fast outlet	1/day	0.001	100
k_1	Parameter of lower slow outlet	1/day	0.001	30
k_{perc}	Percolation parameter	1/day	0.001	30
k_2	Parameter of slow store	1/day	0.001	30

Following parameter optimization, the posterior distribution is approximated using a multistage Metropolis algorithm analogous to the scheme described in *Thyer et al.* [2009]. The modal parameter set from the multistart optimization is used as the initial Metropolis sample. The reported results are based on 20,000 samples, collected following a burn-in stage of 10,000 samples.

For WLS-AR1-PP, the quasi-Newton optimization and MCMC sampling procedures are carried out at each stage of the calibration process (although the hydrological parameter uncertainty estimated in Stage 1 is not used when making predictions in Stage 2, it is useful for model analysis purposes).

3.5. Diagnostics

The predictive performance of the calibration schemes is scrutinized using a set of diagnostics targeted to specific aspects of inferred parameter distributions and predictive distributions (PDs). Note that all diagnostics are applied in the validation period. A diagnostic package based on the RFortran software [*Thyer et al.*, 2011] was used.

Reliability. A predictive distribution is considered to be statistically reliable when it adequately captures the distributional properties of the observed data (i.e., when the observations can be viewed as samples from the PD). In this study, reliability is assessed graphically using predictive quantile-quantile (PQQ) plots [*Laio and Tamea*, 2007; *Thyer et al.*, 2009] and quantified numerically using the metric $\pi_{\text{Reliab}}[\hat{\mathbf{Q}}||\mathbf{Q}]$ presented in Appendix B. This metric is similar to the reliability metric presented in *Renard et al.* [2011] and quantifies the departure of the PQQ plot from the uniform distribution (diagonal PQQ plot). The metric varies between 0 (perfect reliability) and 1 (worst reliability).

Precision. Predictive precision (i.e., the width of the prediction limits) is summarized using the average coefficient of variation (CV) of the predictive distribution

$$CV[\mathbf{Q}] = \frac{1}{N_t} \sum_{t=1}^{N_t} \frac{sdev[Q_t]}{E[Q_t]} \quad (9)$$

where $E[Q_t]$ and $sdev[Q_t]$ denote, respectively, the mean and standard deviation of the predictive streamflow distribution at time t . Both moments are estimated numerically from the MCMC samples. Given multiple PDs with comparable reliability, the distribution with the lowest CV is preferred.

Autocorrelation. The adequacy of the AR(1) assumption in characterizing the error autocorrelation is appraised using the autocorrelation function (ACF) of the innovations $\mathbf{y}$. For a given time series of raw residuals $\mathbf{z}$, the innovations $\mathbf{y}$ are obtained by computing the normalized residuals $\boldsymbol{\eta}$ using equations (4) and (5), followed by solving equation (6) for y_t , $t = 1 \dots N_t$. If the autocorrelation is well captured by the error model, the ACF of the innovations at lags 1 and above, $\rho(\boldsymbol{\eta}_t|\boldsymbol{\eta}_{t-k})$ for $k \geq 1$, should be small (e.g., within the 95% confidence bands of the ACF of white noise).

Heteroscedasticity. The adequacy of the linear heteroscedasticity assumption in equation (5) is appraised by estimating the empirical standard deviations of the innovations, $\sigma_{y(t)}^{(e)}$, as a function of simulated streamflow. This dependence is estimated using a nonparametric kernel technique similar to the approach used by *Villarini and Krajewski* [2008] for rainfall error analysis, except we assume (additive) innovations with a mean of 0, instead of (multiplicative) errors with a mean of 1.

If heteroscedasticity is adequately represented using equation (5), the nonparametric estimates $\sigma_{y(t)}^{(e)}$ will fluctuate around 1 (as we assumed $\sigma_{y(t)} = 1$, see section 2.3). Conversely, systematic deviations from $\sigma_{y(t)}^{(e)} \approx 1$ would suggest that the assumption of linear heteroscedasticity is inadequate.

Statistical properties after aggregation of daily streamflows to a monthly scale. Understanding the predictive performance at aggregated time scales is part of Aim 2 of this study. This is investigated by appraising the predictive distributions of monthly streamflow obtained by aggregating the daily streamflow predictions. The reliability and precision of monthly streamflow PDs are quantified using the same metrics as listed above for the daily flows.

Evaluating parameter inference (synthetic data only). The accuracy of the calibration schemes is appraised by comparing the exact parameter values to the distributions of optimal parameters estimated from individual synthetic replicates.

The reliability of parameter uncertainty estimates is examined using predictive QQ plots. For the k th inferred parameter with exact (synthetic) value $\tilde{\theta}_k$, the following procedure is used:

1. For each replicate of synthetic data, evaluate the cdf of $\tilde{\theta}_k$ within the marginal posterior distribution $p(\theta_k | \tilde{\mathbf{X}}, \tilde{\mathbf{Q}})$. This is repeated for all N synthetic replicates, yielding a set of N cdf values.
2. The distribution of cdf values obtained in step (1) is compared to the uniform distribution using a QQ plot.

Steps (1)–(2) mimic the PQQ approach used to examine predictive streamflow reliability. The parameter inference is said to be reliable if the exact parameters are consistent with being samples from the posterior distributions.

The impact of accounting for error autocorrelation on the parameter uncertainty is investigated using the ratios of standard deviations of marginal posterior distributions. The ratios are computed over the 100 synthetic replicates for each individual parameter inferred using WLS and WLS-AR1.

Evaluating parameter inference (MOPEX catchments). As the exact parameter values are unknown, the inferred parameter distributions are compared between different calibration scenarios.

3.6. Classification of Results Based on Predictive Performance

In order to help summarize and interpret the predictive performance of different calibration schemes in the MOPEX study, the results are grouped according to broad features of predictive performance (see section 5.1). In particular, our aim is to distinguish cases where the predictive distribution is generally precise and reliable from cases with a clear failure to achieve adequate predictions. Albeit based on subjective interpretation, such a classification is helpful for comparing multiple modeling scenarios (12 catchments, 2 hydrological models, 3 calibration schemes, totaling 72 different scenarios).

4. Results of Synthetic Study

4.1. Comparison of Joint and Postprocessor Inference Approaches (Aim 1)

4.1.1. Predictive Performance

Table 4 reports the statistical properties of the predictions for the synthetic experiment. It shows little difference between the performance of the joint (WLS-AR1) and postprocessor (WLS-AR1-PP) approaches, both at the daily and monthly scales. The approximations introduced in the postprocessor approach (section 1.4) have a relatively minor impact on the predictions, at least with synthetic data. Section 5 will examine if this finding holds true for real data.

4.1.2. Parameter Inference

Figure 1 shows the distributions of MAP (Maximum A Posteriori) values for all parameters, estimated in the 100 synthetic runs. The WLS-AR1 and WLS-AR1-PP estimates are generally centered on the “exact” values of the hydrological and error model parameters. An exception is the autocorrelation parameter ϕ , for which the WLS-AR1-PP estimates are slightly lower than expected.

Table 5 summarizes the distributions of the ratios of marginal posterior standard deviations for hydrological and error model parameters inferred using WLS and WLS-AR1. For most parameters, the mean ratios are below 1, indicating that, on average, their posterior uncertainty increases when autocorrelation is included in the error model. A notable exception is the GR4J parameter θ_4 , for which there was on average a reduction in posterior uncertainty.

Table 4. Statistical Properties of the Predictions in the Synthetic Experiment^a

	Daily Scale			Monthly Scale	
	CV[Q]	$\pi_{\text{Reliab}}[\hat{Q} Q]$	$\rho(\eta_t \eta_{t-1})$	CV[Q]	$\pi_{\text{Reliab}}[\hat{Q} Q]$
WLS	0.35 ± 0.01	0.06 ± 0.03	0.88 ± 0.01	0.07 ± 0.002	0.33 ± 0.02
WLS-AR1	0.35 ± 0.01	0.06 ± 0.03	0.01 ± 0.01	0.22 ± 0.01	0.09 ± 0.04
WLS-AR1-PP	0.35 ± 0.01	0.06 ± 0.03	0.01 ± 0.01	0.22 ± 0.01	0.10 ± 0.05

^aThe summary statistics are reported as mean $\pm$ standard deviation, computed for the validation period using 100 replicates.

Figure 1. Distributions of optimal (modal) parameter estimates across the 100 synthetic runs. Vertical red lines indicate the “exact” values of hydrological and error model parameters. The hydrological parameters are not shown for calibration scheme WLS-AR1-PP because it uses the WLS estimates (section 2.3).

4.2. Comparison of Residual Models With and Without Autocorrelation (Aim 2)

4.2.1. Predictive Performance

Table 4 shows little difference in the precision (CV) and reliability (π_{Reliab}) of daily predictions of WLS and WLS-AR1/WLS-AR1-PP. However, the autocorrelation of the innovations is very high for the WLS scheme, because this error model ignores the error persistence.

In clear contrast to the daily results, the monthly performance metrics of WLS are poor compared to WLS-AR1/WLS-AR1-PP. In particular, WLS monthly predictions are statistically unreliable—they significantly underestimate the streamflow uncertainty and produce predictions that are misleadingly “overconfident” with respect to their precision. This will be discussed in section 6.3.

4.2.2. Parameter Inference

Figure 1 shows that the spread (variability) of the optimal estimates of hydrological parameters increases when error autocorrelation is ignored (WLS compared to WLS-AR1). A similar result is obtained for the error heteroscedasticity parameter α_1^* . It can also be seen that the omission of autocorrelation did not introduce parameter biases.

Figure 2 assesses the reliability of the posterior parameter distributions obtained in the 100 synthetic replicates. The WLS posteriors clearly underestimate the uncertainty in the hydrological parameters θ_1 , θ_2 , and θ_3 . Conversely, WLS-AR1 provides very reliable parameter posteriors. Similar results are obtained for parameters α_0^* and α_1^* of the residual error models: WLS posterior distributions underestimate the parameter uncertainty, whereas WLS-AR1 posteriors are very reliable. Thus, ignoring error autocorrelation has a direct impact on the hydrological and error model parameter estimates and leads to posterior distributions that are generally overconfident regarding the parameter uncertainty.

Table 5. Ratios of Marginal Posterior Standard Deviations of GR4J Parameters Estimated Using WLS and WLS-AR1 in the Synthetic Study^a

$r(\theta_1)$	$r(\theta_2)$	$r(\theta_3)$	$r(\theta_4)$
0.44 ± 0.3	0.33 ± 0.02	0.93 ± 0.06	2.59 ± 0.6

^aThe values are reported as mean $\pm$ standard deviation computed over the 100 replicates. The notation $r(\theta)$ denotes $\text{sdev}_{\text{WLS}}[\theta]/\text{sdev}_{\text{WLS-AR1}}[\theta]$, so that ratios below 1 indicate an underestimation of posterior uncertainty when error autocorrelation is ignored.

5. Results of Application to MOPEX Catchments

5.1. Classification of Scenarios Based on Predictive Performance

This section presents the results for 72 scenarios comprising 12 MOPEX catchments, 2 hydrological models, and 3 calibration schemes. Figure 3 shows the precision, reliability and autocorrelation metrics of predictive performance

Figure 2. PQQ plots assessing the reliability of the posterior distributions obtained in the synthetic experiment. Agreement with the 1:1 line indicates that the exact parameter values are consistent with being samples from the posteriors obtained in the 100 individual synthetic calibrations. Gray, white and black symbols correspond to WLS, WLS-AR1, and WLS-AR1-PP, respectively. The symbols are used to distinguish between the curves rather than to indicate individual data points. The hydrological parameters are not shown for calibration scheme WLS-AR1-PP because it uses the optimal (“modal”) WLS values (section 2.3).

for the 72 scenarios. Note that all metrics and diagnostics are reported for the validation period.

To help the presentation and interpretation of the results, a classification scheme based on predictive performance is developed. In particular, we group predictive performance into five classes labeled “good,” “fair,” “mediocre,” “poor,” and “failure.” Figure 4 shows a scatterplot of the precision and reliability metrics for the WLS-AR1 and WLS-AR1-PP scenarios, and provides the framework for classifying the scenarios. Predictive distributions of streamflow typical of each class are presented in Figure 5. The class characteristics are outlined next:

Good class. The predictive distributions have high reliability and precision, with visual inspection showing no obvious prediction problems (see Figure 5, top). Here, we consider high precision to correspond to a CV below 0.4 and high reliability to correspond to π_{Reliab} below 0.1. The fraction of resampled (negative) flows in this class is negligible, below 0.1% (see section 2.4).

Fair, mediocre, and poor classes. These classes describe deterioration in predictive performance, with precision and reliability decreasing from “fair” to “poor” (Figure 5, middle). This corresponds to CVs in the approximate range of 0.4–0.6 and π_{Reliab} increasing from about 0.1 to about 0.4 (Figure 4). The fraction of resampled (negative) flows in these classes ranges from about 1% to 20%.

Failure class. This class represents unacceptable predictive performance arising from a failure of the inference method. It is typified by unrealistically wide prediction limits, with CV in excess of 0.7 (Figure 5, bottom), large fractions of resampled flows (40–50% of flows had to be resampled to avoid negative predictions), along with other evidence of gross failure.

The results in Figure 3 are color coded according to these classes. Figure 6 provides a class-based summary of predictive performance of the joint and postprocessor approaches, for all catchments and hydrological models. Figure 7 reports the diagnostics detailed in section 3.5 for selected representative catchments: the wet catchment B7 (where “good” performance is attained), the cold and snowy catchment B8 (“mediocre” performance), and the dry catchment B12 (“failure” performance).

5.2. Comparison of Joint and Postprocessor Strategies (Aim 1)

5.2.1. Predictive Performance—General

A striking result in Figure 6 is that 5 out of the 12 joint inference scenarios are classed as “failure” for the GR4J model, while no “failures” are noted for the HBV model. Moreover, there are no “failures” in any of the

Figure 3. Metrics of predictive performance at the daily scale. From left to right: CV (precision), π_{Reliab} (reliability), and lag-one autocorrelation (AC) of the innovations. The color codes correspond to the predictive performance classes (see section 5.1 and Figure 4).

postprocessor scenarios. “Good” outcomes are obtained in only one catchment, B7. With the exception of the “failure” cases, there is little variation between the joint and postprocessor approaches—for a given catchment and hydrological model, the predictions of WLS-AR1 and WLS-AR1-PP schemes tend to fall into the same performance class.

A common problem in most scenarios with “mediocre” to “poor” performance is a loss of predictive reliability. This loss is due to significant overprediction, typified by convex-downward PQQ plots (e.g., as seen in row 1 of Figure 7 for catchment B8).

Figure 4. Analysis of predictive performance metrics for hydrological models GR4J and HBV, and calibration schemes WLS-AR1 and WLS-AR1-PP, across all MOPEX catchments. Note that WLS marginal predictions are very similar to WLS-AR1-PP marginal predictions (see end of Appendix A for a brief theoretical discussion). The scenarios classified as “good,” “fair,” “poor,” and “failure” are circled; all other scenarios are classified as “mediocre.” Note that the transition from “fair” to “poor” performance is gradual and hence the classification of these scenarios is clearly subjective. See Figure 5 for an illustration of hydrographs representative of the key classes, and see section 5.1 for further discussion.

Figure 5. Representative streamflow predictions illustrating the performance classification scheme, for catchments B7, B8, and B12. To facilitate the comparison, the secondary y-axis shows the streamflow scaled by the maximum observed flow in that catchment. (top) “Good” class predictions are tight and track the observed data. (middle) “Mediocre” class predictions are wider and provide poorer tracking of observations. (bottom) “Failure” class predictions are unrealistically wide and clearly worse than the other classes (e.g., with some events predicted to have flows up to six times higher than the maximum observed flow, yet with low bounds above observations, etc.).

Catchment	GR4J		HBV	
	WLS-AR1	WLS-AR1-PP	WLS-AR1	WLS-AR1-PP
B1				
B2				
B3				
B4				
B5				
B6				
B7				
B8				
B9				
B10				
B11				
B12				

Good	Fair	Mediocre	Poor	Failure

Figure 6. Classification of predictive performance of GR4J and HBV. Results are shown for the WLS-AR1 and WLS-AR1-PP inference schemes. Except for the five clear failures when GR4J is calibrated using WLS-AR1, the joint and postprocessor approaches have comparable performance.

5.2.2. Predictive Performance—Dependence on the Hydrological Model (Aim 1A)

Figure 6 shows that all five failures of the joint inference approach, namely in catchments B5, B6, B10, B11, and B12, occur when the GR4J hydrological model is used. Inspection of corresponding hydrographs (e.g., Figure 5) and PQQ plots (e.g., Figure 7) indicates that the joint inference approach applied to GR4J produces a major overestimation of predictive uncertainty in combination with an overprediction of streamflow. These defects are not apparent when the HBV hydrological model is used, in which case the joint and postprocessor methods produce a comparable predictive performance and have similar PQQ plots.

5.2.3. Predictive Performance—Dependence on Catchment Properties (Aim 1B)

When GR4J is used, the joint and postprocessor approaches produce comparable, largely mediocre, performance in catchments B1, B2, B4, B7, B8, and B9. The joint approach is only superior in catchment B3. However, in catchments B5, B6, B10, B11, and B12, the postprocessor approach produces “mediocre” to “poor” predictions, while the joint approach fails outright.

When HBV is used, the joint and postprocessor approaches produce comparable, also largely mediocre, performance for most catchments. The exceptions are catchments B3 and B11, where the postprocessor method is superior. Significantly, there are no “failures” for either inference approach.

A cross comparison of the predictive performance in Figure 3 with the catchment characteristics in Table 1 shows that the only catchment with “good” performance is B7, which is the wettest catchment in the MOPEX set. The catchments with “fair” performance are B2 and B3, which have annual rainfall over 1000 mm and runoff coefficients of about 40%. The catchments with “failures”, namely B5, B6, B10, B11, and B12, have no obvious common characteristics. For instance, B11 and B12 have the lowest runoff coefficients and highest ephemerality (>90%), while B5, B6, and B10 have annual rainfall in excess of 1000 mm and ephemerality in the range of 65–70%.

5.2.4. Parameter Inference (Residual Error Models)

This section focuses on error model parameters: the heteroscedasticity parameter α_1^* and the autocorrelation parameter ϕ . These parameters can be compared across all catchments and all hydrological models. The inference of hydrological parameters will be examined in section 5.3.2.

Figure 8 shows that, in general, the estimates of heteroscedasticity and autocorrelation are higher when the joint approach is used. The uncertainty estimates are also higher for the joint strategy. For some scenarios, the posterior distributions of the heteroscedasticity and autocorrelation parameters are particularly different, with the joint approach providing much higher values ($\phi \rightarrow 0.9$, $\alpha_1^* > 1.5$) than the postprocessor strategy. This is much more common with GR4J (catchments B5–B6 and B10–B12) than with HBV. High values of the error model parameters are generally associated with the “failure” class.

Table 6 presents the correlations between the GR4J groundwater exchange parameter θ_2 , the heteroscedasticity parameter α_1 , and the autocorrelation parameter ϕ , in individual catchments, for calibration scheme WLS-AR1. Strong correlations are present between α_1 and θ_2 . In 7 out of 12 catchments, the correlation is

Figure 7. Diagnostics of predictive performance of WLS, WLS-AR1, and WLS-AR1-PP, for representative catchments B7 (“good”), B8 (“mediocre”), and B12 (“failure”). (row 1; Reliability) Predictive QQ plots. (row 2; Heteroscedasticity) Standard deviations of innovations versus quantiles of simulated flow. (row 3; Distributional check) Probability density plots of innovations, with the black line showing a standard Gaussian pdf. (row 4; Autocorrelation) ACF of the innovations, with 95% confidence intervals indicated with dotted lines. In rows 1–3, the symbols are used to distinguish between the curves rather than to indicate individual data points.

stronger than 0.9, with the five “failure” catchments having even stronger correlations of 0.99 (catchments B5, B6, B10, and B12) and 0.96 (catchment B11). However, the correlations of ϕ versus θ_2 and ϕ versus α_1 are low, mostly in the range of 0.2–0.4.

The interactions between these parameters are further examined in Figure 9, which plots the modal estimates of θ_2 versus α_1^* and θ_2 versus ϕ across all catchments. It shows strong associations between these pairs of parameters. The “failure” catchments have highly negative values of θ_2 (indicating significant groundwater export) and, concomitantly, high values of α_1^* and ϕ . As shall be discussed in section 6.1, the combination of these factors can lead to a spurious inflation of predictive uncertainty.

Figure 8. Posterior distributions of the marginal heteroscedasticity parameter α_1^* and the error autocorrelation parameter ϕ , color coded according to the performance classes. The red boxes on the α_1^* axis boundary indicate values exceeding 1.5.

5.3. Comparing Residual Models With and Without Autocorrelation (Aim 2)

5.3.1. Predictive Performance

Figure 3 shows that WLS and WLS-AR1-PP predictions have very similar reliability and precision. However, the autocorrelation is considerably better represented when the residual error model contains the AR(1) component. Comparing WLS and WLS-AR1, the precision is generally worse with WLS-AR1, especially when the GR4J hydrological model is used. For a few scenarios, the reliability improves for WLS-AR1 compared to WLS (B1 and B4 for GR4J, B6 for HBV) at the expense of a slight loss of precision.

Density plots of the estimated innovations (row 3 of Figure 7) reveal that they are asymmetric and biased for WLS (except for catchment B7, where they are approximately Gaussian). The estimated innovations are symmetric but kurtotic for the WLS-AR1 and WLS-AR1-PP schemes. The kurtosis is stronger in catchments B8 and B11. The Gaussian assumptions are hence questionable in these cases.

Table 6. Correlations (Pearson Coefficient) Between the GR4J Water Balance Parameter θ_2 , the Error Heteroscedasticity Parameter α_1 , and the Error Autocorrelation Parameter ϕ , Obtained in the 12 MOPEX Catchments^a

Error Model	WLS-AR1			WLS	WLS-AR1-PP
Catchment	$\rho(\theta_2, \phi)$	$\rho(\theta_2, \alpha_1)$	$\rho(\phi, \alpha_1)$	Stage 1 $\rho(\theta_2, \alpha_1)$	Stage 2 $\rho(\phi, \alpha_1)$
B1	−0.09	−0.59	0.18	−0.47	0.01
B2	−0.27	−0.96	0.30	−0.48	0.01
B3	−0.02	−0.44	0.02	−0.46	0.00
B4	−0.23	−0.73	0.18	−0.60	0.00
B5	−0.59	−0.99	0.59	−0.50	0.00
B6	−0.07	−0.99	0.07	−0.48	−0.02
B7	0.04	−0.77	0.14	−0.31	0.01
B8	−0.16	−0.85	0.24	−0.64	0.02
B9	−0.25	−0.91	0.27	−0.63	−0.02
B10	−0.35	−0.99	0.35	−0.51	0.00
B11	−0.24	−0.96	0.25	−0.44	0.02
B12	−0.35	−0.99	0.35	−0.29	0.03

^aCorrelations with magnitude exceeding 0.9 are indicated in **bold** font.

When autocorrelation is ignored in the residual error model (calibration scheme WLS), large autocorrelations are present in the actual innovations (right column, Figure 3 and row 4 of Figure 7). In calibration schemes WLS-AR1 and WLS-AR1-PP, the AR(1) assumption adequately reproduces the autocorrelation structure of the normalized residuals, with the ACF of the actual innovations being close to 0 in all three catchments.

At the monthly time scale, aggregated WLS predictions are generally too precise (overconfident) and unreliable, as seen in Figure 10. The predictive uncertainty of monthly streamflow is thus strongly underestimated by the WLS approach. This is attributed to neglecting the strong autocorrelation of residual errors. The use of WLS-AR1 and WLS-AR1-PP schemes significantly rectifies this shortcoming. This is further discussed in section 6.3.

5.3.2. Parameter Inference

Figure 11 shows examples of the parameter uncertainty for catchments B7, B8, and B11 using the WLS and WLS-AR1 schemes. Comparing the uncertainty of the hydrological model parameters between WLS and WLS-AR1, it can be seen that including autocorrelation in the error model leads to markedly different parameter estimates and posterior distributions. There is no clear trend of increased parameter uncertainty when the autocorrelation is taken into account, with results varying for different catchments and hydrological models. For example, when catchment B7 is modeled using GR4J, the posterior uncertainty increases for parameters θ_1 , θ_2 , and θ_3 , but decreases for parameter θ_4 . These results are similar to the synthetic case

Figure 9. Relationship (a) between modal (most probable) WLS-AR1 estimates of GR4J groundwater exchange parameter θ_2 and the marginal heteroscedasticity parameter α_1^* and (b) between θ_2 and the error autocorrelation parameter ϕ . Results for all 12 MOPEX catchments are shown, color coded according to predictive performance. The estimates obtained using WLS (i.e., ignoring error autocorrelation) are shown for reference in Figure 9a. Note that all “failure” scenarios have clearly elevated values of α_1^* and ϕ , and are associated with very large negative values of θ_2 .

Figure 10. Predictive distributions of aggregated monthly streamflow for catchment B7 and GR4J model. WLS predictions, which ignore the error autocorrelation at the daily scale, significantly underestimate predictive uncertainty at the monthly scale.

study (section 4 and Table 6). In contrast, for HBV, when autocorrelation is included, parameter uncertainty decreases for six of the eight parameters (S_{FC} , β , γ_{PWP} , L , k_0 , and k_1) yet increases for two parameters (k_{perk} and k_2). This finding illustrates that, though in most cases the inclusion of error autocorrelation results in considerably different parameter estimates, the precise effect is difficult to predict a priori. This issue is further discussed in section 6.4.

6. Discussion

6.1. When Inferring Error Autocorrelation, the Postprocessor Approach Can Be More Robust Than the Joint Approach

Despite the theoretical attraction of the joint inference approach, the empirical results suggest that joint inference of hydrological parameters and the residual error autocorrelation parameter can be nonrobust. In particular, the joint approach can produce unrealistically wide predictive limits.

The nonrobustness appears related to a water balance parameter in the GR4J model. In particular, the GR4J groundwater exchange parameter θ_2 , the error heteroscedasticity parameter α_1 , and the error autocorrelation parameter ϕ can interact to produce exceedingly wide uncertainty in the predicted streamflows (e.g., Figure 5, catchment B12). This overestimation of predictive uncertainty is spurious because it does not occur when the same model is calibrated using WLS or WLS-AR1-PP schemes.

The problematic behavior of the joint inference arises from subtle parameter interactions. When calibrating GR4J using WLS-AR1 in a given catchment, the correlations of ϕ versus θ_2 and ϕ versus α_1 are low (Table 6). However, when considering these relationships over all catchments a different picture emerges. As shown in Figure 9, in “failure” catchments, θ_2 tends toward large negative values (corresponding to massive groundwater export), α_1 increases to anomalously high values (1 and above), and ϕ approaches 1 which is the limit of stationarity of the AR(1) process. The combination of high values of α_1 and ϕ produce the extremely wide prediction limits that lead to performance failure.

Conversely, there is no evidence of WLS-AR1 failing when calibrating the HBV model. For example, Figure 3 shows substantially lower CV values when HBV is applied to the same catchments where GR4J failed. On

Figure 11. Comparison of the posterior distributions of the hydrological parameters obtained using the residual error models WLS and WLS-AR1 in catchments B7, B8, and B12. Including autocorrelation in the error model produces clear differences in the calibrated parameter values and uncertainty estimates

the basis of the above, the absence of WLS-AR1 failures with HBV could be tentatively attributed to the absence of a water balance parameter analogous to θ_2 .

However, synthetic testing indicates that this degeneracy does not occur when WLS-AR1 is used to calibrate GR4J to synthetic data with substantial water export ($\theta_2 = -5$ mm/d) and highly heteroscedastic and autocorrelated errors ($\alpha_1 = 0.33$ and $\phi = 0.92$) (results not shown). The absence of failure under synthetic conditions suggests that the nonrobustness of the joint approach with real data may not be due solely to the presence of water balance parameters. Structural errors in the hydrological and/or error models and observational errors in the data are also likely contributory factors.

More generally, most hydrological models have parameters that control the overall water balance. For example, both GR4J and HBV have parameters that control the actual evapotranspiration (ET) as a fraction of potential ET and hence affect the total streamflow volume predicted by the model. However, the effect of the groundwater exchange parameter in GR4J is much stronger than the effect of typical ET parameters: (i) unlike the actual ET, which is constrained by the potential ET, groundwater exports in GR4J are essentially only constrained by how much water is available in the model storage; (ii) the groundwater flux occurs

much latter in the flow simulation chain than ET and thus has a more direct effect on the final streamflow; (iii) the groundwater flux in GR4J is implemented using a fixed-step explicit Euler time stepping scheme, which is susceptible to numerical artefacts [Kavetski and Clark, 2010]; and (iv) GR4J has a relatively low number of calibrated parameters, which may restrict its predictive ability once the water balance parameter is affected by problematic interactions. Some or all of these factors may explain why, compared to HBV, GR4J appears vulnerable to parameter interactions within the joint inference approach.

Further research, including detailed experimentation with different hydrological model structures, parameterizations, and prior information, is required to understand and ameliorate the problematic parameter interactions within the joint inference approach. In the meantime, the postprocessor approach appears to be a more robust practical choice. This finding does not necessarily preclude the use of the joint approach as a “method of first choice,” but does highlight that its results may be nonrobust and require recomputation.

6.2. Why is the Postprocessor Approach More Robust Despite Its Apparent Theoretical Shortcomings?

The major theoretical shortcoming of the postprocessor approach is that it ignores interactions between hydrological and residual error model parameters. However, it is precisely the omission of these parameter interactions that makes the postprocessor approach more robust in practice.

Another theoretical shortcoming of the postprocessor approach used in this study is that the streamflow predictions are conditioned on the optimal hydrological parameter set and therefore ignore the uncertainty in the hydrological parameters. However, under the assumption that data and structural errors are described using the residual error model, and given a long calibration period, posterior parametric uncertainty is low. As a result, predictive uncertainty is dominated by the uncertainty associated with the residual error model.

In addition to producing reliable and precise streamflow predictions, an important goal of an inference framework is to produce model parameter estimates interpretable in terms of hydrological processes in particular catchments. Due to the simplified error model used when calibrating the hydrological model and the conditional calibration of the residual error model, parameter estimates obtained using postprocessor approaches may be unreliable, both in terms of point estimates and the associated uncertainties. From a strict theoretical perspective, posterior parameter distributions obtained using the joint inference are hence preferable to postprocessor estimates, as seen in the synthetic study. However, in practice, the nonrobustness of the joint approach also undermines its parameter estimates, e.g., spuriously large values of the groundwater export parameter shown in Figure 9. Overall, the practical advantages provided by the postprocessor approach for streamflow prediction do not negate the need to further improve the error models underlying the likelihood function.

6.3. When Does Incorporating Error Autocorrelation Improve Predictive Performance?

At the daily time scale, the reliability and precision metrics show little difference between the performance of WLS and WLS-AR1-PP. This indicates that including autocorrelation in the residual error model makes little difference to the reliability and precision of marginal predictions at the daily time scale (the marginal daily predictions obtained using WLS and WLS-AR1-PP are very similar; see end of Appendix A for a brief theoretical discussion). However, analysis of innovations (Figure 7, row 4) highlights the temporal structure for the residual errors and suggests that the inclusion of the AR(1) model is both justified and necessary.

The first visible advantage of including autocorrelation in the residual error model is the improved “realism” of individual streamflow time series replicates. For example, Figure 12a compares streamflow time series simulated using WLS versus WLS-AR1. The independent random nature of the WLS error model produces “noisy” streamflow time series. In contrast, the WLS-AR1 error model provides a “smoother” time series that is qualitatively more similar to actual observed streamflow.

The second advantage of including error autocorrelation is in forecasting contexts, where newly available observations can be used to enhance future predictions (e.g., Montanari and Grossi [2008] and others). Figure 12b compares one-day-ahead conditional predictive distributions obtained for representative streamflow events using the WLS and WLS-AR1 error models (Appendix A). WLS-AR1 predictions take advantage of streamflow observations from the preceding time step, producing clearly better precision and tracking of observations than WLS. For example, in Figure 12b the conditional prediction using WLS-AR1 “dynamically”

Figure 12. Representative events illustrating the value of including autocorrelation in the error model. Results for catchment B7 and hydrological model GR4J are shown. (a) Individual stochastic replicates of streamflow generated using WLS versus WLS-AR1 error models (same random number sequence was used when sampling the residual errors). Autocorrelated errors (WLS-AR1) are smoother due to their persistence and produce more realistic streamflow time series compared to time series with independent errors (WLS), where the streamflow is too “noisy.” (b) One-day ahead conditional probabilistic predictions (Appendix A). Methods that include error autocorrelation (WLS-AR1) take advantage of preceding observations, providing higher precision and better tracking of next-step data.

corrects the hydrological model to better track the recession curve. As residual errors of hydrological models are typically highly persistent (e.g., Figure 9 shows autocorrelation values in the range of 0.7 - 0.98 for the MOPEX catchments), using autocorrelated error models will generally be of major benefit for conditional forecasting.

The third practical advantage of including error autocorrelation is the improvement in the reliability of temporally aggregated flows. Figure 10 clearly demonstrates that when flows are aggregated to the monthly time scale, WLS significantly underestimates the predictive uncertainty and provides predictive limits that are grossly overconfident. The underestimation of uncertainty can be explained by considering that the variance of a time series aggregated over n_{agg} time steps is the sum of the variances and covariances of the original time series,

$$\text{var} \left[\sum_{i=1}^{n_{agg}} Q_i \right] = \sum_{i=1}^{n_{agg}} \text{var} [Q_i] + 2 \sum_{i=1}^{n_{agg}} \sum_{j=1}^{i-1} \text{covar} [Q_i, Q_j] \quad (10)$$

If the covariance terms are ignored, as in a residual error model with no autocorrelation, the variance of the aggregated time series will be underestimated.

This variance aggregation behavior has important practical implications for operational contexts such as “seamless” streamflow forecasting using rainfall-runoff models, where streamflow predictions are produced

at one time scale (e.g., daily) and are then aggregated to longer time scales (e.g., weekly, monthly, and seasonal) for practical use.

6.4. Incorporating Autocorrelation in the Residual Errors Does Not Always Increase the Estimated Parameter Uncertainty

When errors are autocorrelated, the effective sample size is generally reduced [Brillinger, 1989; Hamed and Rao, 1998]. Hence, accounting for error autocorrelation recognizes the lower information content of the data. In turn, at least intuitively, this should lead to higher estimates of parameter uncertainty compared to approaches that disregard error autocorrelation.

This expectation is largely confirmed by the synthetic study, where the inclusion of autocorrelation increases uncertainty in six out of seven parameters (Table 6). The exception, for parameter θ_4 , suggests that including autocorrelation does not always result in increased parameter uncertainty.

While the WLS optimal parameter estimates are unbiased, their increased spread (Figure 1) indicates that WLS parameter estimates are on average less accurate than WLS-AR1. This corresponds to a loss of statistical efficiency compared to the correct error model, WLS-AR1. This is accompanied by an underestimation of posterior parameter uncertainty (Figure 2), implying the estimates are unreliable.

For the MOPEX data, the results show that including autocorrelation in the residual error model produces substantial changes in estimated parameter uncertainty. The specific behavior varies considerably between catchments and hydrological models. For some parameters, the estimated uncertainty markedly increases, while for others it decreases. Inspection of Figure 11 suggests that the reduction in uncertainty occurred primarily for parameters that control the recession, but this requires further work to confirm and explain more conclusively.

There are also parameters for which switching from WLS to WLS-AR1 produces significant shifts in the estimates themselves. This behavior occurs for both hydrological and error model parameters. The lack of replication in the real data study precludes statements regarding whether the shifts in parameter estimates correspond to genuine estimation biases or merely to sampling variability. Further work is needed to better understand the reasons for these changes in the parameter estimates.

6.5. Dependence of Predictive Performance on Catchment Characteristics

In catchments B2, B7, and, to a lesser extent B3, “good” to “fair” predictive performance is obtained irrespective of the hydrological model and calibration strategy. For example, for catchment B7, reliability (Figure 7) and treatment of heteroscedasticity (Figure 7) are adequate for all scenarios. The common feature of catchments B2, B3, and B7 is that they have relatively high runoff coefficients, and rank relatively low in terms of the degree of ephemerality (Table 1). However, these characteristics are not sufficient to ensure good performance. For example, despite having runoff coefficients and ephemerality very similar to catchment B2, catchments B5, B6, and B10 have mediocre-poor performance (and catchment B10 fails outright when GR4J is calibrated using WLS-AR1).

Catchments B8, B9, B11, and B12 exhibit mediocre/failed performance irrespective of the calibration strategy or hydrological model used. A common characteristic of these catchments is their comparatively low runoff coefficient and high ephemerality (Table 1). These ephemeral catchments are challenging for the hydrological models used in this study. For example, neither GR4J nor HBV have thresholds in their storage-discharge relationships, making it difficult to reproduce “spiky” hydrograph behavior. In turn, this leads to more complex error characteristics, e.g., (i) heteroscedasticity does not follow the linear form of equation (5); (ii) dispersion of residual errors tends to be overestimated for low flows and underestimated for high flows in catchment B8 (Figure 7); and (iii) a systematic underestimation of dispersion is observed for catchment B11 when HBV is applied. Snow dynamics (catchments B1–B8) add another aspect of hydrological complexity, in particular the timing of snow melt, which may be responsible for timing errors in the predicted hydrographs.

7. Conclusions and Recommendations

7.1. Joint Versus Postprocessor Approaches for Inferring Error Autocorrelation

When autocorrelation is included in the residual error model, the postprocessor approach (WLS-AR1-PP scheme) appears more robust than the joint inference approach (WLS-AR1 scheme). The joint approach is

clearly nonrobust for the calibration of the GR4J hydrological model, due to strong interactions between the GR4J groundwater exchange parameter, the error heteroscedasticity parameter, and the error autocorrelation parameter. Therefore, care must be taken when inferring error autocorrelation jointly with hydrological parameters, especially those that directly affect water balance. If strong parameter interactions arise, a postprocessor approach may provide more robust parameter estimates and predictions (at the expense of foregoing joint inference of all parameters). Consequently, despite some theoretical reservations discussed in section 6.2, the postprocessor approach can currently be recommended as the more robust option for practical applications focusing on streamflow prediction.

When comparing predictive performance across catchments, the best results are obtained for the French Broad River, which is the wettest and least ephemeral catchment, while the worst results are generally obtained in catchments that are highly ephemeral. With the exception of the joint approach failures for the GR4J model, the joint and postprocessor approaches produce comparable and largely mediocre predictive performance. We generally attribute low predictive precision to structural errors in the hydrological model and low predictive reliability to inadequate characterization of residual errors, particularly at low flows.

7.2. Importance of Incorporating Error Autocorrelation

A major practical value of incorporating autocorrelation in the errors of hydrological predictions is the clear improvement in predictive uncertainty quantification at aggregated time scales. Conversely, the uncertainty in predicted streamflows is strongly underestimated when the error persistence is not taken into account. For example, when the hydrological model is calibrated at a daily scale using the WLS calibration scheme, the predicted streamflows aggregated to a monthly scale substantially underestimate predictive uncertainty. Furthermore, accounting for error autocorrelation produces smoother and visibly more “realistic” simulated streamflow time series at the daily time scale.

Incorporating autocorrelation also has a significant impact on the parameter uncertainty estimation. However, the results vary from catchment to catchment: in some cases the uncertainty is reduced, in other cases it increases, and in others the parameter distributions are significantly shifted.

The results also suggest that the inclusion of error autocorrelation makes little difference on the overall (marginal) reliability and precision of streamflow predictions that are not conditioned on immediate past observations. However, in forecasting applications, knowledge of error autocorrelation can be exploited to significantly improve the predictions by taking advantage of newly arriving observations.

7.3. Future Research

The inclusion of a wide range of catchments in this study shows that some catchments have poor predictive performance irrespective of the hydrological model and/or calibration approach. These problematic catchments tend to have a high degree of ephemerality and/or significant snow processes. To improve parameter estimation and predictive performance in these types of catchment, both hydrological models and residual errors models require further development.

In terms of future development of residual errors models with autocorrelation, the following directions are of interest: (i) understanding and, if possible, remedying problematic multiway interactions between hydrological parameter(s) and residual error model parameters; (ii) the presence of bias in the model predictions (e.g., due to model structural error) requires careful consideration and treatment [Reichert and Schuwirth, 2012]; these tasks are challenging because there could be major interactions of water balance parameters in the hydrological model with bias and autocorrelation parameters in the residual error model; (iii) the treatment of frequent zero and near-zero flows in ephemeral catchments requires further attention, including approaches that specify a discrete probability mass at zero streamflow [Smith *et al.*, 2010]; (iv) the error heteroscedasticity in ephemeral catchments is poorly represented by a linear function of simulated flows [Evin *et al.*, 2013]; as the present study used a linear representation of heteroscedasticity, the development and analysis of nonlinear heteroscedastic error models and transformational approaches is recommended.

Appendix A: Predictive Uncertainty Estimation Using the WLS-AR1 Error Model

Predictive distributions corresponding to the WLS-AR1 error model are obtained from equations (1–6).

A.1 - Marginal Predictions

Marginal predictive distributions are used when streamflow observations at preceding time steps are not available or are not used for conditioning. Given a parameter set $(\theta_H, \theta_\varepsilon)$, the k th sample from the marginal distribution of predicted streamflow time series, $\mathbf{Q}^{(k)} \equiv \mathbf{Q}_{1:t}^{(k)}$, can be computed as follows:

$$\eta_0 = 0; y_1^{(k)} \leftarrow N(0, 1/(1-\phi^2)) \quad (\text{A1})$$

$$\eta_{t+1}^{(k)} = \phi \eta_t^{(k)} + y_{t+1}^{(k)}; y_{t+1}^{(k)} \leftarrow N(0, 1^2) \quad (\text{A2})$$

$$\varepsilon_t^{(k)} = \sigma_t \eta_t^{(k)}; \sigma_t = \alpha_0 + \alpha_1 Q_t^{\theta_H} \forall t \quad (\text{A3})$$

$$\mathbf{Q}^{(k)} = \mathbf{Q}^{\theta_H} + \boldsymbol{\varepsilon}^{(k)} \quad (\text{A4})$$

where $Q_t^{\theta_H}$ is the streamflow at time t computed using the hydrological model with parameter set θ_H as defined in equation (1). Note that $\mathbf{Q}^{(k)}$ corresponds to a sampled "prediction" of observed streamflow and is directly comparable to the actual observed values $\tilde{\mathbf{Q}}$ (as discussed at the end of section 2.3).

The specification of $(\theta_H, \theta_\varepsilon)$ depends on the calibration scheme. In the joint approach, equations (A2)–(A4) are iterated over $(\theta_H, \theta_\varepsilon)^{(i)}$ where i indexes samples from the full posterior $p(\theta_H, \theta_\varepsilon | \tilde{\mathbf{X}}, \tilde{\mathbf{Q}})$, whereas the post-processor approach iterates over $(\theta_H^{(WLS)}, \theta_\varepsilon^{(j)})$ where j indexes samples from the conditional posterior $p(\theta_\varepsilon | \theta_H^{(WLS)}, \tilde{\mathbf{X}}, \tilde{\mathbf{Q}})$. Additional variations of the postprocessor approach can be constructed using alternative specifications of $(\theta_H, \theta_\varepsilon)$, but are not considered in this work.

A.2 - Conditional Predictions

Consider one-step-ahead prediction, i.e., the prediction of streamflow Q_{t+1} given observed streamflow $\tilde{Q}_t$. More formally, we are interested in the conditional predictive distribution $p(Q_{t+1} | \tilde{Q}_t)$. In this case, equation (A2) can be applied without sampling from the marginal distribution of η_t and instead directly computing η_t from the observation $\tilde{Q}_t$. Substituting equations (1) and (4) into equation (A2), and making use of $\tilde{Q}_t$ to compute $\varepsilon_t = \tilde{Q}_t - Q_t^{\theta_H}$, yields

$$\frac{Q_{t+1}^{(k)} - Q_{t+1}^{\theta_H}}{\sigma_{t+1}} = \phi \frac{\tilde{Q}_t - Q_t^{\theta_H}}{\sigma_t} + y_{t+1}^{(k)}; y_{t+1}^{(k)} \leftarrow N(0, 1^2) \quad (\text{A5})$$

Equation (A5) can be solved for $Q_{t+1}^{(k)}$ to yield

$$Q_{t+1}^{(k)} = Q_{t+1}^{\theta_H} + \phi \frac{\sigma_{t+1}}{\sigma_t} (\tilde{Q}_t - Q_t^{\theta_H}) + \sigma_{t+1} y_{t+1}^{(k)}; y_{t+1}^{(k)} \leftarrow N(0, 1^2) \quad (\text{A6})$$

and hence the one-step ahead predictive distribution can be sampled as follows:

$$Q_{t+1}^{(k)} | \tilde{Q}_t \leftarrow N\left(Q_{t+1}^{\theta_H} + \phi \frac{\sigma_{t+1}}{\sigma_t} (\tilde{Q}_t - Q_t^{\theta_H}), \sigma_{t+1}^2\right) \quad (\text{A7})$$

The conditioning on $\tilde{Q}_t$ provides two distinct advantages: (i) the mean of the predictions is shifted to account for the persistence of the errors, generally reducing predictive bias and improving reliability, and (ii) the conditional variance of the predictions is smaller than the marginal variance (especially when ϕ is near 1 as seen in equation (8)), generally yielding much more precise (tighter) predictions.

Finally, conditional prediction can be generalized to n -step-ahead prediction. As n increases, the conditional predictions converge to the marginal predictions. In particular, n -step-ahead WLS-AR1-PP predictions converge to the WLS predictions because the schemes differ solely in the inclusion of autocorrelation in the residual error model used for predictions and therefore have equivalent marginal predictive distributions. The marginal WLS-AR1 predictions will generally be similar to the marginal WLS and WLS-AR1-PP predictions, though they will not be identical due to potential differences in hydrological parameter values (in some calibrations, such as those in the "failure" class discussed in section 5.1, the parameter differences between WLS-AR1 and WLS might be substantial and result in markedly different predictions).

Appendix B: Reliability Metric

A predictive distribution (PD) is considered to be statistically reliable when it adequately captures the distributional properties of the observed data, i.e., when the observations can be viewed as samples from the PD. Reliability can be assessed graphically using predictive quantile-quantile (PQQ) plots [Laio and Tamea, 2007; Thyer et al., 2009]. A PQQ plot displays the empirical cumulative distribution function (cdf) of the set Ω , which is defined as the set of cdf values of observed data (e.g., streamflow) within the predictive distribution. More formally,

$$\Omega[\tilde{\mathbf{Q}}|\mathbf{Q}] = \{F_{Q(t)}(\tilde{Q}_t), t=1..N_t\} \quad (\text{B1})$$

where $\tilde{Q}_t$ is the observation at time t and $F_{Q(t)}(\cdot)$ is the predictive cdf at time step t .

If the observed data are consistent with being samples from the predictive distribution, the elements of Ω will be uniformly distributed. Therefore, a quantitative reliability metric can be derived by considering a norm of the difference between the PQQ plot and the uniform distribution (diagonal cdf),

$$\pi_{\text{Reliab}}[\tilde{\mathbf{Q}}|\mathbf{Q}] = \|F_U - F_{\Omega[\mathbf{Q}|\tilde{\mathbf{Q}}]}^{(e)}\| \quad (\text{B2})$$

where F_U is the uniform cdf and $F_{\Omega}^{(e)}(\cdot)$ is the empirical cdf of the set of samples Ω .

As a specific choice of norm in equation (B2), we use the 1-norm, i.e., the sum of absolute discrepancies at each observed cdf value. The resulting metric can be written in full as

$$\pi_{\text{Reliab}}[\tilde{\mathbf{Q}}|\mathbf{Q}] = \frac{2}{N_t} \sum_{t=1}^{N_t} \left| F_U[F_{Q(t)}(\tilde{Q}_t)] - F_{\Omega[\mathbf{Q}|\tilde{\mathbf{Q}}]}^{(e)}[F_{Q(t)}(\tilde{Q}_t)] \right| \quad (\text{B3})$$

where the factor $2/N_t$ is included to ensure the metric lies between 0 (perfect reliability) and 1 (worst reliability).

The reliability metric in equation (B3) is similar to the metric used by Renard et al. [2011]. It quantifies predictive reliability by the extent to which the observed values of a quantity of interest (here, streamflow time series) are uniformly distributed within the time series of predictive distributions (which is expected if $\tilde{Q}_t$ is consistent with being a sample from $F_{Q(t)}$ at all time steps t).

Acknowledgments

This work was carried as part of Australian Research Council grant DP110105408. The MOPEX data set is available on request from Qingyun Duan. We thank Ben Renard for useful comments on an earlier draft of the manuscript. We thank the editors and reviewers for constructive comments and suggestions during the review of the manuscript.

References

- Bates, B. C., and E. P. Campbell (2001), A Markov chain Monte Carlo scheme for parameter estimation and inference in conceptual rainfall-runoff modeling, *Water Resour. Res.*, 37(4), 937–947, doi:10.1029/2000WR900363.
- Bergström, S. (1995), The HBV model, in *Computer Models of Watershed Hydrology*, edited by V. P. Singh, pp. 443–476, Water Resour. Publ., Highlands Ranch, Colo.
- Beven, K., P. Smith, I. Westerberg, and J. Freer (2012), Comment on “Pursuing the method of multiple working hypotheses for hydrological modeling” by P. Clark et al., *Water Resour. Res.*, 48, W11801, doi:10.1029/2012WR012282.
- Beven, K. J., and A. M. Binley (1992), The future of distributed hydrological models: Model calibration and uncertainty prediction, *Hydrol. Processes*, 6(3), 279–298, doi:10.1002/hyp.3360060305.
- Beven, K. J., P. J. Smith, and J. E. Freer (2008), So just why would a modeller choose to be incoherent?, *J. Hydrol.*, 354(1–4), 15–32, doi:10.1016/j.jhydrol.2008.02.007.
- Brillinger, D. R. (1989), Consistent detection of a monotonic trend superposed on a stationary time series, *Biometrika*, 76, 23–30.
- Clark, M. P., and D. Kavetski (2010), Ancient numerical daemons of conceptual hydrological modeling. Part 1: Fidelity and efficiency of time stepping schemes, *Water Resour. Res.*, 46, W10510, doi:10.1029/2009WR008894.
- Clark, M. P., A. G. Slater, D. E. Rupp, R. A. Woods, J. A. Vrugt, H. V. Gupta, T. Wagener, and L. E. Hay (2008), Framework for Understanding Structural Errors (FUSE): A modular framework to diagnose differences between hydrological models, *Water Resour. Res.*, 44, W00B02, doi:10.1029/2007WR006735.
- Clark, M. P., D. Kavetski, and F. Fenicia (2012), Reply to comment by K. Beven et al. on “Pursuing the method of multiple working hypotheses for hydrological modeling”, *Water Resour. Res.*, 48, W11802, doi:10.1029/2012WR012547.
- Duan, Q., et al. (2006), Model Parameter Estimation Experiment (MOPEX): An overview of science strategy and major results from the second and third workshops, *J. Hydrol.*, 320(1–2), 3–17.
- Engeland, K., B. Renard, I. Steinsland, and S. Kolberg (2010), Evaluation of statistical models for forecast errors from the HBV model, *J. Hydrol.*, 384(1–2), 142–155, doi:10.1016/j.jhydrol.2010.01.018.
- Evin, G., D. Kavetski, M. Thyer, and G. Kuczera (2013), Pitfalls and improvements in the joint inference of heteroscedasticity and autocorrelation in hydrological model calibration, *Water Resour. Res.*, 49, 4518–4524, doi:10.1002/wrcr.20284.
- Hamed, K. H., and A. R. Rao (1998), A modified Mann-Kendall trend test for autocorrelated data, *J. Hydrol.*, 204(1–4), 182–196.

- Kavetski, D., and M. P. Clark (2010), Ancient numerical daemons of conceptual hydrological modeling. Part 2: Impact of time stepping scheme on model analysis and prediction, *Water Resour. Res.*, **46**, W10511, doi:10.1029/2009WR008896.
- Kavetski, D., S. W. Franks, and G. Kuczera (2003), Confronting input uncertainty in environmental modelling, in *Calibration of Watershed Models*, edited, pp. 49–68, AGU, Washington, D. C.
- Kuczera, G. (1983), Improved parameter inference in catchment models: 1. Evaluating parameter uncertainty, *Water Resour. Res.*, **19**(5), 1151–1162, doi:10.1029/WR019i005p01151.
- Laio, F., and S. Tamea (2007), Verification tools for probabilistic forecasts of continuous hydrological variables, *Hydrol. Earth Syst. Sci.*, **11**(4), 1267–1277, doi:10.5194/hessd-3-2145-2006.
- Mantovan, P., and E. Todini (2006), Hydrological forecasting uncertainty assessment: Incoherence of the GLUE methodology, *J. Hydrol.*, **330**(1–2), 368–381, doi:10.1016/j.jhydrol.2006.04.046.
- Montanari, A. (2007), What do we mean by 'uncertainty'? The need for a consistent wording about uncertainty assessment in hydrology, *Hydrol. Processes*, **21**, 841–847.
- Montanari, A., and G. Grossi (2008), Estimating the uncertainty of hydrological forecasts: A statistical approach, *Water Resour. Res.*, **44**, W00B08, doi:10.1029/2008WR006897.
- Montanari, A., and D. Koutsoyiannis (2012), A blueprint for process-based modeling of uncertain hydrological systems, *Water Resour. Res.*, **48**, W09555, doi:10.1029/2011WR011412.
- Morawietz, M., C.-Y. Xu, L. Gottschalk, and L. M. Tallaksen (2011), Systematic evaluation of autoregressive error models as post-processors for a probabilistic streamflow forecast system, *J. Hydrol.*, **407**(1–4), 58–72, doi:10.1016/j.jhydrol.2011.07.007.
- Perrin, C., C. Michel, and V. Andreassian (2001), Does a large number of parameters enhance model performance? Comparative assessment of common catchment model structures on 429 catchments, *J. Hydrol.*, **242**(3–4), 275–301, doi:10.1016/S0022-1694(00)00393-0.
- Perrin, C., C. Michel, and V. Andreassian (2003), Improvement of a parsimonious model for streamflow simulation, *J. Hydrol.*, **279**(1–4), 275–289, doi:10.1016/S0022-1694(03)00225-7.
- Pianosi, F., and L. Raso (2012), Dynamic modeling of predictive uncertainty by regression on absolute errors, *Water Resour. Res.*, **48**, W03516, doi:10.1029/2011WR010603.
- Reichert, P., and J. Mieleitner (2009), Analyzing input and structural uncertainty of nonlinear dynamic models with stochastic, time-dependent parameters, *Water Resour. Res.*, **45**, W10402, doi:10.1029/2009WR007814.
- Reichert, P., and N. Schuwirth (2012), Linking statistical bias description to multiobjective model calibration, *Water Resour. Res.*, **48**, W09543, doi:10.1029/2011WR011391.
- Renard, B., D. Kavetski, M. Thyer, G. Kuczera, and S. W. Franks (2010), Understanding predictive uncertainty in hydrologic modeling: The challenge of identifying input and structural errors, *Water Resour. Res.*, **46**, W05521, doi:10.1029/2009WR008328.
- Renard, B., D. Kavetski, E. Leblois, M. Thyer, G. Kuczera, and S. W. Franks (2011), Toward a reliable decomposition of predictive uncertainty in hydrological modeling: Characterizing rainfall errors using conditional simulation, *Water Resour. Res.*, **47**, W11516, doi:10.1029/2011WR010643.
- Schaeffli, B., D. B. Talamba, and A. Musy (2007), Quantifying hydrological modeling errors through a mixture of normal distributions, *J. Hydrol.*, **332**(3–4), 303–315, doi:10.1016/j.jhydrol.2006.07.005.
- Schoups, G., and J. A. Vrugt (2010), A formal likelihood function for parameter and predictive inference of hydrologic models with correlated, heteroscedastic and non-Gaussian errors, *Water Resour. Res.*, **46**, W10531, doi:10.1029/2009WR008933.
- Seibert, J. (1997), *HBV Light Version 1.2, User's Manual*, Uppsala Univ., Dep. of Earth Sci., Hydrol., Uppsala, Sweden.
- Smith, T., A. Sharma, L. Marshall, R. Mehrotra, and S. Sisson (2010), Development of a formal likelihood function for improved Bayesian inference of ephemeral catchments, *Water Resour. Res.*, **46**, W12551, doi:10.1029/2010WR009514.
- Sorooshian, S., and J. A. Dracup (1980), Stochastic parameter estimation procedures for hydrological rainfall-runoff models: Correlated and heteroscedastic error cases, *Water Resour. Res.*, **16**(2), 430–442.
- Stedinger, J. R., D. P. Lettenmaier, and R. Vogel (1985), ARMA(1,1) Models for multiple site annual streamflow synthesis, *Water Resour. Res.*, **21**(4), 497–510.
- Stedinger, J. R., R. M. Vogel, S. U. Lee, and R. Batchelder (2008), Appraisal of the Generalized Likelihood Uncertainty Estimation (GLUE) method, *Water Resour. Res.*, **44**, W00B06, doi:10.1029/2008WR006833.
- Thyer, M., B. Renard, D. Kavetski, G. Kuczera, S. Franks, and S. Srikanthan (2009), Critical evaluation of parameter consistency and predictive uncertainty in hydrological modelling: A case study using bayesian total error analysis, *Water Resour. Res.*, **45**, W00B14, doi:10.1029/2008WR006825.
- Thyer, M., M. Leonard, D. Kavetski, S. Need, and B. Renard (2011), The open source RFortran library for accessing R from Fortran, with applications in environmental modelling, *Environ. Modell. Software*, **26**(2), 219–234, doi:10.1016/j.envsoft.2010.05.007.
- Tuteja, N. K., et al. (2011), *Experimental Evaluation of the Dynamic Seasonal Streamflow Forecasting Approach*, Bur. of Meteorol., Melbourne, Victoria.
- van Andel, S. J., A. Weerts, J. Schaake, and K. Bogner (2013), Post-processing hydrological ensemble predictions intercomparison experiment, *Hydrol. Processes*, **27**(1), 158–161, doi:10.1002/hyp.9595.
- Villarini, G., and W. F. Krajewski (2008), Empirically based modeling of spatial sampling uncertainties associated with rainfall measurements by rain gauges, *J. Hydrol.*, **31**(7), 1015–1023, doi:10.1016/j.advwatres.2008.04.007.
- Vrugt, J. A., C. G. H. Diks, H. V. Gupta, W. Bouten, and J. M. Verstraten (2005), Improved treatment of uncertainty in hydrologic modeling: Combining the strengths of global optimization and data assimilation, *Water Resour. Res.*, **41**, W01017, doi:10.1029/2004WR003059.
- Ye, W., B. C. Bates, N. R. Viney, M. Sivapalan, and A. J. Jakeman (1997), Performance of conceptual rainfall-runoff models in low-yielding ephemeral catchments, *Water Resour. Res.*, **33**(1), 153–166, doi:10.1029/96WR02840.
- Zhang, Z., T. Wagener, P. Reed, and R. Bhushan (2008), Reducing uncertainty in predictions in ungauged basins by combining hydrologic indices regionalization and multiobjective optimization, *Water Resour. Res.*, **44**, W00B04, doi:10.1029/2008WR006833.