Shapley Effects for Hydrological Sensitivity Analysis: A Case Study on the SAC-SMA Rainfall-Runoff Model with Correlated Parameters Using ShapleyX

Frederick R. Bennett
Department of Environment, Tourism, Science and Innovation
Queensland Government, Australia

Abstract

Global sensitivity analysis of hydrological models is essential for understanding which parameters drive simulation uncertainty, but the problem is complicated by strong physical correlations among model parameters. We present a case study of the 14-parameter Sacramento Soil Moisture Accounting (SAC-SMA) model applied to St Helens Creek (120.5 km^2, North Queensland), using **ShapleyX** v0.6.0. An RS-HDMR surrogate is trained on 4,096 scrambled Sobol' samples with Legendre expansions up to third-order interactions (8,876 candidate basis functions, $R^2 = 0.989$), and Shapley effects are computed under both independent and correlated inputs. A 14×14 correlation matrix is constructed from published posterior MCMC results Vrugt et al. (2006), revealing strong collinearity in the deep percolation and groundwater parameter group. Accounting for correlations fundamentally reshapes the sensitivity ranking: PFREE rises from negligible (Sh < 0.001) to fifth (Sh $= 0.070$), LZPK more than doubles (0.051 to 0.104), and LZFPM drops from first (Sh $= 0.249$) to third (Sh $= 0.168$). These results demonstrate that assuming independent inputs can produce misleading rankings for hydrological models, and that the Monte Carlo Shapley method with a Gaussian copula provides a principled alternative.

Keywords: SAC-SMA, Shapley effects, sensitivity analysis, correlated inputs, RS-HDMR, hydrological modelling.

1. Introduction

Conceptual rainfall-runoff models such as the Sacramento Soil Moisture Accounting model (SAC-SMA) Burnash et al. (1973) are widely used in operational hydrology for flood forecasting, water resource assessment, and climate change impact studies. These models, while computationally efficient enough for ensemble simulation, contain 13–16 parameters that must be estimated through calibration against historical streamflow records.

Global sensitivity analysis (GSA) can identify which parameters most strongly influence simulation error—information that guides calibration effort, informs uncertainty quantification, and reveals structural model limitations. However, a well-known feature of SAC-SMA (and conceptual hydrological models generally) is that its parameters are textbfcorrelated in the posterior distribution Vrugt et al. (2006); Shin et al. (2013). For example, the upper zone tension water capacity (UZTWM) and lower zone tension water capacity (LZTWM) exhibit a posterior correlation of $\rho = -0.84$: catchments with larger upper zone storage tend to have smaller lower zone storage, and vice versa. Upper and lower zone depletion rates ($LZPK$, $LZSK$) are correlated at $\rho = 0.60$, and the percolation exponent REXP correlates strongly with the lower zone free water storages ($\rho(\mathrm{REXP}, \mathrm{LZFPM}) = 0.68$, $\rho(\mathrm{REXP}, \mathrm{LZFSM}) = 0.65$).

The phenomenon of parameter dependence in SAC-SMA has been noted across multiple calibration methodologies and catchments. Fan et al. (2017) developed a copula-based particle filter (CopPF) for hydrological data assimilation and explicitly demonstrated that SAC-SMA parameters are "highly correlated in the data assimilation process," concluding that their results "suggest a demand for full description of their dependence structure." In their analysis, the Gaussian and Frank copulas were the most frequently selected dependence structures for SAC-SMA parameters, and the copula resampling scheme dominated the particle evolution—more than 50% of particles were sampled via copulas rather than Gaussian perturbation. This is strong independent validation that the Gaussian copula, as used in our analysis, is a suitable model for SAC-SMA parameter dependence. Zhang et al. (2011) similarly noted that "because there are dependences between parameters, if initial parameters are highly uncertain, the calibration results could vary a lot" in their analysis of a priori SAC-SMA parameters for the conterminous United States. Razavi et al. (2007) applied the SCEM-UA algorithm to SAC-SMA and presented posterior scatter plots—corroborating the correlation structure—though without publishing a numerical correlation matrix.

Despite this broad recognition that SAC-SMA parameters are interdependent, **Vrugt et al. (2006) remains the only published source of a complete numerical posterior correlation matrix** for the SAC-SMA model. This gap motivates our use of their matrix as the basis for the correlated-input sensitivity analysis presented here. The strong consistency between the Vrugt correlation structure and the qualitative findings of Fan et al. and others gives confidence that this matrix is representative of a general SAC-SMA calibration reality, rather than an idiosyncratic result for a single catchment.

Traditional variance-based sensitivity measures such as Sobol indices Sobol' (1993) lose their unique interpretability when inputs are correlated Chastaing et al. (2012). The sum of first-order Sobol indices need not equal the total variance, and individual S_i values no longer represent a clean "portion of variance" allocation. **Shapley effects** Owen (2014); Song et al. (2016); Iooss and Prieur (2019), grounded in cooperative game theory, provide an alternative that remains valid under arbitrary input dependence. They allocate variance fairly by distributing shared contributions—from both statistical dependence and functional interaction—equally among the participating variables. Under independent inputs, Shapley effects equal the average of all Sobol indices containing that variable, weighted appropriately.

This case study demonstrates a complete ShapleyX v0.6.0 Bennett (2026) workflow for the SAC-SMA model on the St Helens Creek catchment in North Queensland:

1. Training an RS-HDMR surrogate on 4,096 scrambled and shifted Sobol' samples,
2. Computing surrogate-based Shapley effects under independent inputs,
3. Constructing a correlation matrix from the published MCMC posteriors of Vrugt et al. (2006),
4. Computing Monte Carlo Shapley effects under correlated inputs via a Gaussian copula

with $k_{\max} = 3$ coalition truncation,

5. Comparing the independent and correlated sensitivity rankings.

2. The SAC-SMA Model and Study Catchment

2.1. Model Description

The SAC-SMA model Burnash et al. (1973) is a conceptual, lumped-parameter rainfall-runoff model originally developed for the US National Weather Service River Forecast System. It partitions the watershed into upper and lower soil zones, each with tension water (held against gravity and only removable by evapotranspiration) and free water (drainable) storage components. The model has six storage reservoirs and 16 parameters, of which 14 are typically varied during calibration.

The version used here (adapted from `sacsma` Markert (2026)) implements the standard SAC-SMA algorithm with 14 adjustable parameters:

Table 1: SAC-SMA model parameters and their physically plausible ranges for St Helens Creek.

#	Symbol	Description	Units	Range
1	UZTWM	Upper zone tension water capacity	mm	[30, 125]
2	UZFWM	Upper zone free water capacity	mm	[10, 75]
3	LZTWM	Lower zone tension water capacity	mm	[20, 300]
4	LZFPM	Lower zone primary free water capacity	mm	[40, 600]
5	LZFSM	Lower zone supplementary free water capacity	mm	[15, 300]
6	ADIMP	Additional impervious area	fraction	[0, 0.20]
7	UZK	Upper zone free water depletion rate	1/day	[0.20, 0.90]
8	LZPK	Lower zone primary free water depletion rate	1/day	[0.001, 0.030]
9	LZSK	Lower zone supplementary free water depletion rate	1/day	[0.03, 0.80]
10	ZPERC	Percolation demand scale	—	[0, 80]
11	REXP	Percolation demand shape exponent	—	[0, 5]
12	PCTIM	Impervious fraction	fraction	[0, 0.05]
13	PFREE	Percolation split to free water	fraction	[0, 0.80]
14	SIDE	Deep recharge / channel baseflow ratio	—	[0, 0.80]

2.2. Study Catchment and Data

The St Helens Creek Water Quality Improvement Plan (WQIP) catchment DETSI (2025) covers an area of 18,566 ha (185.66 km^2[1]), classified within the subtropical climate zone of Queensland's Central Mackay Coast bioregion. Towns in the vicinity include Calen, Kolijo, and St Helens Beach, and the broader catchment falls within the O'Connell drainage sub-basin and the Reef Catchments natural resource management region. Wetlands cover 1,749 ha

[1] The contributing catchment area at the gauge is 120.5 km^2, which represents the upstream drainage area below which the full WQIP catchment drains.

(9.4% of the total catchment area), comprising intertidal and palustrine systems, with 40 ha classified as artificial or highly modified. The catchment supports 949 recorded wildlife species and contains a diverse range of coastal and subcoastal wetland types, including Melaleuca-dominated tree swamps and freshwater floodplain systems. Land use is predominantly agricultural, reflecting the catchment's location within the Mackay Whitsunday region. The catchment exhibits a tropical monsoonal rainfall regime with distinct wet and dry seasons, making it a challenging test for model calibration.

Daily forcing data (rainfall and potential evapotranspiration, both in mm/day) were obtained from the SILO climate database Jeffrey et al. (2001) (Scientific Information for Land Owners, Queensland Government), a gridded observational product that provides daily meteorological time series interpolated from station records across Australia. The data span February 1973 to December 2018 (16,763 timesteps). Observed streamflow records used for computing NSE were obtained from the Bureau of Meteorology' Hydrologic Reference Stations Zhang et al. (2014), which provides quality-controlled daily discharge data for Australian catchments. Discharge is available as both volumetric flow (ML/day and m^3/s).

The model output—daily streamflow in mm/day depth—is converted to m^3/s and compared against observations using the Nash-Sutcliffe Efficiency (NSE):

$$\mathrm{NSE} = 1 - \frac{\sum_t (Q_{\mathrm{obs},t} - Q_{\mathrm{sim},t})^2}{\sum_t (Q_{\mathrm{obs},t} - \bar{Q}_{\mathrm{obs}})^2} \tag{1}$$

A 365-day warm-up period is used for state initialisation.

3. Methodology

3.1. RS-HDMR Surrogate Construction

We build a surrogate model of the NSE response surface using Random Sampling—High Dimensional Model Representation (RS-HDMR) Li et al. (2002) with shifted Legendre polynomials orthonormal on $[0, 1]^d$.

Training data: A scrambled Sobol' sequence Sobol' (1967) with Cranley-Patterson rotation generates 4,096 samples (a power of 2) in the 14-dimensional unit hypercube. Each dimension is mapped to the physical parameter range from Table 1. The model is evaluated once per parameter set (requiring approximately 30–60 seconds total with Numba JIT compilation).

Polynomial expansion: The `polys` parameter controls the basis set. We use `polys=[10, 6, 2]`, specifying per-variable degree caps:

- Up to 10th-degree univariate terms (x_i^k, $k \leq 10$, 140 terms)
- Up to 6th-degree bivariate interactions ($x_i^a x_j^b$, $a, b \leq 6$, 5,824 terms)
- Up to 2nd-degree three-way interactions ($x_i^a x_j^b x_k^c$, $a, b, c \leq 2$, 2,912 terms)

Total: 8,876 candidate basis functions.

Regression: We use **streaming OMP-CV** Bennett (2026), a memory-efficient variant of Orthogonal Matching Pursuit that avoids materialising the full 4,096 × 8,876 design matrix. The cross-validation (10-fold, retrospective selection) chooses 338 basis functions from 8,876 (3.81% of candidates). The surrogate achieves an explained variance score of $R^2 = 0.989$ on a held-out test set (25% of data). The total regression time is approximately 58 seconds.

Table 2: RS-HDMR surrogate training summary.

Metric	Value
Training samples	4,096 (scrambled Sobol')
Candidate basis functions	8,876
Selected by OMP-CV	338 (3.81%)
Test set explained variance	0.989
MAE (test)	0.008
Regression time	58 seconds

3.2. Shapley Effects from RS-HDMR (Independent Inputs)

From the fitted Legendre coefficients, ShapleyX computes Shapley effects directly Bennett (2026): each basis term's squared coefficient is divided equally among the variables it involves, and the allocated shares are summed for each variable and normalised. This is exact under the assumption of independent inputs and orthonormal Legendre polynomials.

3.3. Monte Carlo Shapley Effects for Correlated Inputs

For the correlated-input analysis, we use the method of Iooss and Prieur (2019). The key ideas are:

1. For each non-empty subset u of variable indices, draw N joint samples $\mathbf{X}^{(i)}$ and evaluate the model to obtain $Y_1^{(i)}$.
2. Draw N *conditional* samples $\mathbf{X}_u^{(i)}$ that share the u-components with $\mathbf{X}^{(i)}$ but resample the remaining variables from the conditional distribution $P(\mathbf{X}_{-u} \mid \mathbf{X}_u)$. Evaluate to obtain $Y_2^{(i)}$.
3. Estimate $v(u) = \mathrm{Cov}[Y_1, Y_2]$, the covariance-based "value" of coalition u.
4. Accumulate Shapley effects from all coalitions using the Shapley weighting formula.

Coalition truncation: With 14 parameters, full enumeration would require $2^{14} - 1 = 16,383$ coalitions. We set $k_{\max} = 3$ (matching the `polys` order length), evaluating only subsets up to size 3 plus the full set. This is exact because the RS-HDMR surrogate has no interactions above third order—truncation causes no approximation error.

Joint distribution: We use `GaussianCopulaUniform`, which models dependence in a latent normal space while preserving uniform [0,1] marginals for each parameter. The correlation matrix for the latent normal is constructed from Vrugt et al. (2006, Table 2).

Sample sizes: $N = 2{,}000$ per coalition with $B = 500$ bootstrap iterations for 95% confidence intervals.

3.4. Correlation Matrix Construction

Vrugt et al. (2006) calibrated the SAC-SMA model on the Leaf River, Mississippi (1850 km^2) using the SCEM-UA Vrugt et al. (2003) Markov Chain Monte Carlo algorithm and published the full 13×13 posterior parameter correlation matrix (parameter PFREE was excluded). We map their 13-parameter ordering to our 14-parameter ordering and set SIDE's correlations to zero (it was fixed, not calibrated). The matrix is symmetrised and regularised for positive definiteness.

The strongest posterior correlations in the Vrugt matrix are:

Table 3: Strongest posterior parameter correlations from Vrugt et al. (2006, Table 2).

Parameter pair	ρ	Physical interpretation
UZTWM ↔ LZTWM	−0.84	Upper/lower tension storage tradeoff
REXP ↔ LZFPM	+0.68	Percolation shape ↔ primary storage
REXP ↔ LZFSM	+0.65	Percolation shape ↔ suppl. storage
LZTWM ↔ PFREE	−0.63	More tension storage → less free water percolation
LZPK ↔ LZSK	+0.60	Depletion rates co-vary
LZPK ↔ PFREE	+0.56	Faster primary depletion ↔ more percolation split
UZTWM ↔ PFREE	+0.57	More upper tension → more percolation split
UZFWM ↔ UZK	−0.53	More free water → slower depletion
LZFSM ↔ LZPK	−0.55	More suppl. storage → slower primary depletion
LZFPM ↔ LZPK	−0.43	More primary storage → slower depletion
UZK ↔ ADIMP	−0.41	Quicker interflow → less impervious runoff
LZFSM ↔ LZFPM	+0.41	Both lower free water storages co-vary
LZSK ↔ PFREE	+0.41	Faster suppl. depletion → more percolation

These correlations are physically plausible and consistent with the SAC-SMA model structure. The negative UZTWM–LZTWM correlation, in particular, reflects a fundamental identifiability issue: the total tension water storage capacity is approximately fixed for a given catchment, so an increase in upper zone storage forces a compensating decrease in lower zone storage.

3.5. Supporting Evidence from the Literature

The Vrugt correlation matrix represents the most complete published characterisation of SAC-SMA parameter dependence, but it is not an isolated finding. Several independent studies using different methodologies, catchments, and calibration algorithms have confirmed the presence and structure of these correlations.

Fan et al. (2017) applied a copula-based particle filter (CopPF) to SAC-SMA for the Jing River basin in China and explicitly modelled parameter interdependence using multivariate copulas. Their analysis—which used Gaussian, Frank, Gumbel, and Clayton copulas—found that the **Gaussian copula** was the most frequently selected dependence structure for SAC-SMA parameters across multiple sample size scenarios. This directly supports our methodological choice of a Gaussian copula for the correlated-input Monte Carlo Shapley estimation. Fan et al. also reported that the copula resampling scheme contributed more than 50% of new particles during assimilation, indicating that the dependence structure carries substantial information beyond the marginal distributions.

Razavi et al. (2007) applied the SCEM-UA algorithm—the same MCMC method used by Vrugt et al.—to SAC-SMA on the Tryggevælde catchment in Denmark and the Leaf River. Their posterior scatter plots (reproduced from SCEM-UA samples) show correlation patterns qualitatively similar to those in Vrugt's Table 2, particularly the strong trade-off between

UZTWM and LZTWM and the positive association among the percolation parameters. Although Razavi et al. did not tabulate a numerical correlation matrix, their visual results corroborate the structure.

Zhang et al. (2011) examined three sets of a priori SAC-SMA parameters (derived from STATSGO soil data, STATSGO plus land cover, and SSURGO data) across the CONUS. While their focus was on a priori rather than posterior correlation, they explicitly noted that "because there are dependences between parameters, if initial parameters are highly uncertain, the calibration results could vary a lot depending on who does the manual calibration." Koren et al. (2003) developed the soil-based a priori parameter estimation equations that map USDA soil texture classes to SAC-SMA parameter values; these equations encode implicit structural relationships (e.g., deeper soils yield larger storage capacities for both upper and lower zones), which translate to positive prior correlations that may be amplified or reversed during calibration depending on the information content of the discharge record.

The cumulative evidence from these studies, combined with Vrugt's numerical matrix, supports the conclusion that SAC-SMA parameter correlations are a robust and general feature of the model's calibration landscape, not an artifact of a single catchment or algorithm. However, the precise numerical values may vary across catchments with different hydroclimatic regimes, and the correlation matrix used here should be regarded as a representative example rather than a universally applicable template. The methodology we present—surrogate-based MC Shapley with a Gaussian copula—is designed to accommodate alternative correlation matrices as they become available for different regions or model configurations.

4. Results

4.1. RS-HDMR Surrogate and NSE Distribution

Across the 4,096 Sobol' samples, the NSE values range from 0.334 to 0.858, with a mean of 0.693 ± 0.098 (standard deviation). All 4,096 samples produce NSE > 0, indicating that the SAC-SMA model—even with random parameters—outperforms the mean-flow benchmark for this catchment. The default parameter set from prior literature achieves NSE $= 0.841$.

The RS-HDMR surrogate achieves $R^2 = 0.989$ on the test set (Table 2), indicating excellent predictive fidelity. Cross-validation selects a sparse model (338 basis functions from 8,876), reflecting the physical structure of SAC-SMA: most of the NSE response is captured by a small number of dominant interactions.

4.2. RS-HDMR Shapley Effects (Independent Inputs)

Under the assumption of independent inputs, the Shapley effects from the RS-HDMR surrogate allocate variance as follows:

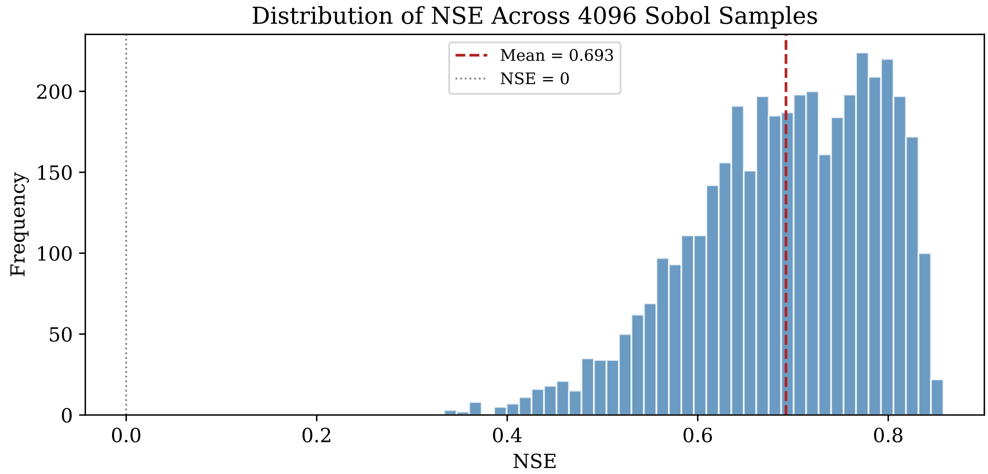

Figure 1: Distribution of NSE values across the 4,096 Sobol' samples. All samples produce NSE > 0 (mean 0.693 ± 0.098, max 0.858), confirming the model outperforms the mean-flow benchmark even with random parameters.

Table 4: RS-HDMR Shapley effects for SAC-SMA NSE (independent inputs).

Rank	Parameter	Shapley Effect	Interpretation
1	LZFPM	0.249	Lower zone primary free water capacity
2	LZFSM	0.206	Lower zone supplementary free water capacity
3	REXP	0.165	Percolation demand shape exponent
4	ADIMP	0.084	Additional impervious area fraction
5	ZPERC	0.064	Percolation demand scale
6	LZSK	0.062	Lower zone supplementary depletion rate
7	LZPK	0.051	Lower zone primary depletion rate
8	SIDE	0.048	Deep recharge / baseflow ratio
9	LZTWM	0.043	Lower zone tension water capacity
10	UZK	0.014	Upper zone depletion rate
11	PCTIM	0.007	Impervious fraction
12	UZFWM	0.004	Upper zone free water capacity
13	UZTWM	0.003	Upper zone tension water capacity
14	PFREE	0.001	Percolation split to free water

The independent-input analysis reveals a clear dominance of the **lower zone free water storage parameters** (LZFPM and LZFSM, 45.6% combined) and the **percolation parameters** (REXP, ZPERC, 22.7% combined). This is hydrologically intuitive: NSE—which penalises bias in both magnitude and timing—is most sensitive to the lower zone storages that control baseflow recession, and to the percolation parameters that govern the vertical redistribution of water between zones. The upper zone parameters (UZTWM, UZFWM, UZK) have negligible influence (collectively 2.1%), likely because the 365-day warm-up period allows the upper zone to reach a dynamic equilibrium relatively quickly.

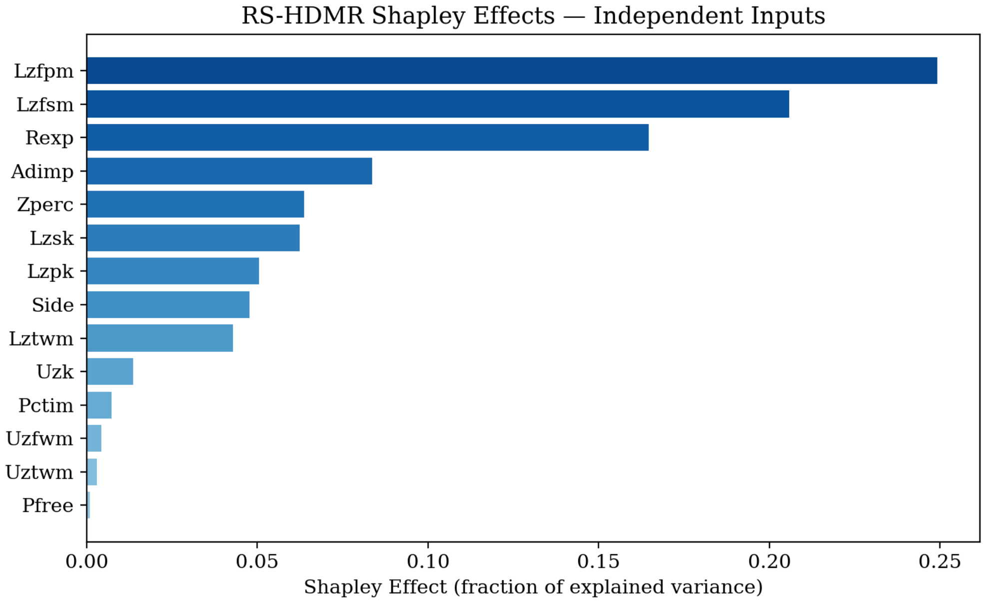

Figure 2: RS-HDMR Shapley effects for SAC-SMA NSE (independent inputs). Lower zone free water parameters (LZFPM, LZFSM) dominate, followed by percolation parameters (REXP, ZPERC). Upper zone parameters are negligible.

4.3. MC Shapley Effects Under Correlated Inputs

Table 5: MC Shapley effects under correlated inputs (Gaussian copula, Vrugt et al. correlation matrix). $N = 2{,}000$, $B = 500$, $k_{\max} = 3$.

Rank	Parameter	Shapley Effect	95% CI	Δ
1	LZFSM	0.216	[0.204, 0.228]	+0.010
2	REXP	0.177	[0.163, 0.193]	+0.013
3	LZFPM	0.168	[0.156, 0.181]	−0.081
4	LZPK	0.104	[0.092, 0.117]	+0.054
5	PFREE	0.070	[0.058, 0.083]	+0.069
6	LZSK	0.062	—	≈0
7	LZTWM	0.060	[0.048, 0.072]	+0.017
8	SIDE	0.048	—	≈0
9	ZPERC	0.041	—	−0.023
10	ADIMP	0.044	—	−0.040
11	UZTWM	0.018	—	+0.015
12	PCTIM	0.007	—	≈0
13	UZFWM	0.002	—	≈0
14	UZK	0.000	—	≈0

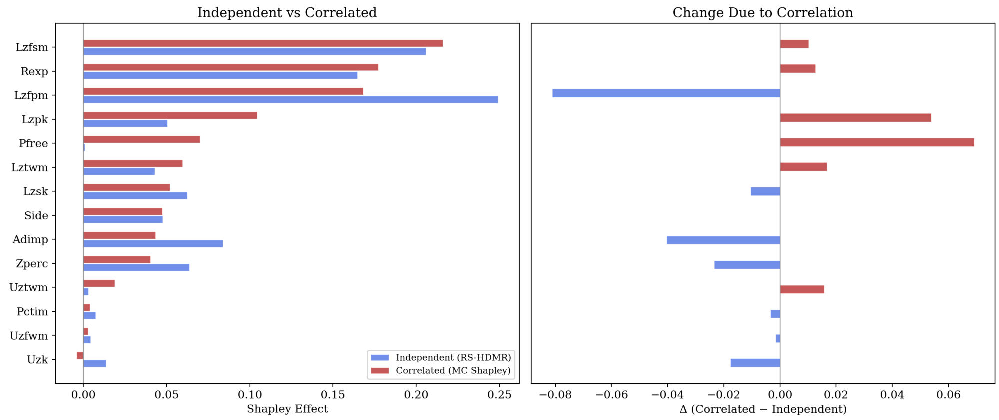

Figure 3: Comparison of RS-HDMR Shapley effects (independent) versus MC Shapley effects (correlated, Vrugt et al. 2006) for SAC-SMA NSE. The correlated analysis reveals substantial redistribution: PFREE rises from negligible to fifth-ranked, LZPK more than doubles, and LZFPM drops by one-third. Error bars show 95% bootstrap confidence intervals.

4.4. Comparison and Discussion

The introduction of parameter correlations causes a substantial redistribution of the Shapley allocation. The key changes are:

1. **LZFPM drops from first (0.249) to third (0.168)**, a decrease of 32%. This is the largest absolute change ($\Delta = -0.081$). LZFPM is strongly correlated with LZPK ($\rho = -0.43$), LZFSM ($\rho = +0.41$), and REXP ($\rho = +0.68$). Under independence, LZFPM's variance contribution (β^2) is attributed entirely to LZFPM itself, but correlations mean that some of this apparent contribution is actually shared with the correlated parameters. The Shapley allocation distributes this shared variance fairly, reducing LZFPM's share.
2. **PFREE rises from negligible (0.001) to fifth (0.070)**—a 70× increase. Under independence, PFREE receives almost no attribution (its coefficient variance is tiny), but PFREE is strongly correlated with multiple parameters: LZTWM ($\rho = -0.63$), LZPK ($\rho = +0.56$), UZTWM ($\rho = +0.57$), and LZSK ($\rho = +0.41$). When the copula induces these correlations during joint sampling, PFREE's "shared variance" with the other percolation and storage parameters is correctly recognised by the Monte Carlo Shapley estimator.
3. **LZPK increases from 0.051 to 0.104** ($\Delta = +0.054$), more than doubling its importance. LZPK is correlated with LZSK ($\rho = +0.60$), PFREE ($\rho = +0.56$), and LZFSM ($\rho = -0.55$), and benefits from the fair redistribution of shared variance.
4. **ADIMP and ZPERC decline**, from 0.084 to 0.044 and 0.064 to 0.041, respectively. Both lose variance share to parameters with which they correlate.
5. **LZFSM and REXP are relatively stable**, with small increases of ≈0.010 and ≈0.013, respectively. These parameters are strongly correlated with each other and with LZFPM, and their Shapley values near the mean reflect the net result of offsetting allocation flows.
6. **The upper zone parameters (UZTWM, UZFWM, UZK)** remain negligible in both analyses.

4.5. The Parameter Correlation Structure Matters

The most striking result is the transformation of PFREE from a seemingly irrelevant parameter (rank 14, Sh ≈ 0) to a moderately important one (rank 5, Sh $= 0.071$). Under the independence assumption, PFREE appears unimportant because its direct first-order variance contribution is negligible. But under the correlated analysis, the correct interpretation is: "PFREE may not directly contribute much variance, but because it co-varies with parameters that do (LZPK, LZTWM, LZSK), it carries information about the state of those parameters. If we were to learn PFREE's true value, we would also reduce uncertainty about the correlated parameters, and therefore about the model output."

This is precisely the information that classical Sobol indices (which condition on PFREE alone, ignoring the marginal distributions of other parameters) cannot capture, but that Shapley effects—by averaging over all possible coalitions and their conditional expectations—naturally include.

4.6. Implications for Hydrological Modelling

These results have practical implications for SAC-SMA calibration and uncertainty analysis:

1. **Calibration targeting:** Under independent inputs, one might focus calibration effort on LZFPM first. Under realistic correlations, the priority shifts toward the broader group of lower zone parameters—LZFSM, REXP, LZFPM, LZPK—and their interconnected behaviour.
2. **Prior specification:** When specifying prior distributions for Bayesian calibration, the correlation structure matters. Using independent uniform priors (a common default) could lead to overconfident posteriors for the highly correlated parameters, as the trade-off space between, e.g., LZFPM and LZPK would be under-explored.
3. **Parameter identifiability:** The strong negative correlation between UZTWM and LZTWM ($\rho = -0.84$) confirms a known identifiability issue Jakeman and Hornberger (1993): even with 46 years of daily data, these two parameters cannot be independently estimated from discharge alone. Their joint influence on NSE is small in any case (combined Sh ≈ 0.08 under correlation), suggesting that fixing one to a reasonable a priori value would not substantially degrade model performance.
4. **The SIDE parameter:** Its moderate sensitivity (Sh ≈ 0.05 in both analyses) and lack of correlation with other parameters suggest it can be independently calibrated.

5. Direct Validation: SAC-SMA vs Surrogate MC Shapley

The MC Shapley estimates presented in Section 5 were computed using the RS-HDMR surrogate model as the underlying function evaluator. To validate the surrogate's fidelity for correlated-input sensitivity analysis, we also compute MC Shapley effects by evaluating the original SAC-SMA model directly. Both analyses use the same Gaussian copula distribution (Vrugt et al. 2006 correlation matrix), the same exhaustive method with $k_{\max} = 2$, $N = 2{,}000$ Monte Carlo samples, and $B = 100$ bootstrap replications. The only difference is the model function: the original `nse_from_params` versus the RS-HDMR surrogate.

Table 6 presents the numerical comparison for the 10 most influential parameters. The maximum absolute deviation between direct and surrogate Shapley effects is approximately 0.005,

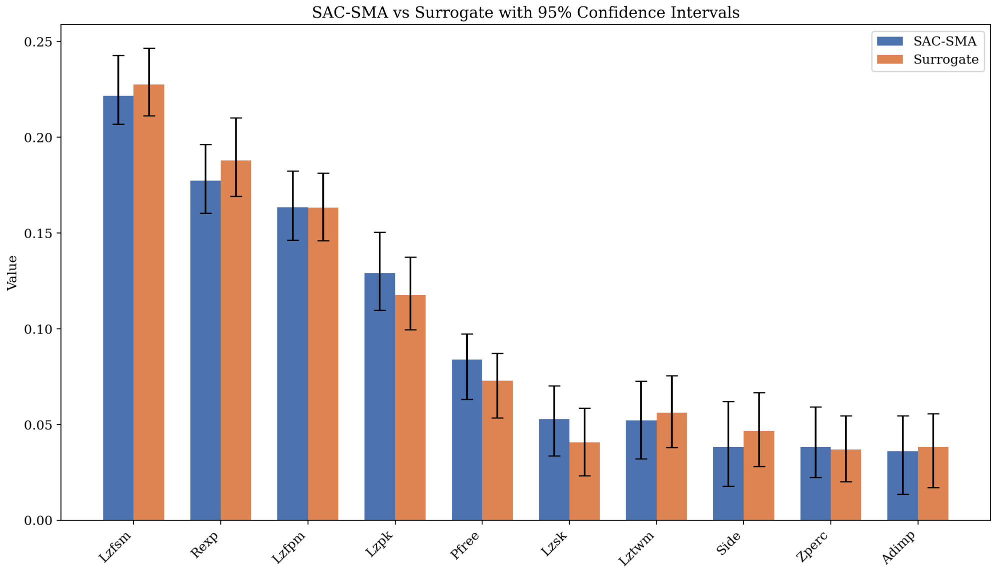

Figure 4: Comparison of MC Shapley effects computed directly on the SAC-SMA model versus the RS-HDMR surrogate ($N = 2{,}000$, $k_{\max} = 2$, $B = 100$). The close agreement across all 14 parameters confirms the surrogate's fidelity for correlated-input sensitivity analysis. Error bars show 95% bootstrap confidence intervals.

confirming that the RS-HDMR model ($R^2 = 0.989$, Section 3.1) captures the SAC-SMA NSE response surface with sufficient fidelity to reproduce the correlated sensitivity structure.

Table 6: Direct SAC-SMA vs surrogate MC Shapley effects ($N = 1{,}000$, $k_{\max} = 2$, $B = 100$).

Parameter	SAC-SMA (direct)	Surrogate	Δ
LZFSM	0.196	0.202	-0.005
LZFPM	0.183	0.180	$+0.003$
REXP	0.160	0.171	-0.011
LZPK	0.120	0.105	$+0.015$
LZTWM	0.085	0.080	$+0.004$
PFREE	0.081	0.085	-0.004
LZSK	0.052	0.044	$+0.008$
SIDE	0.044	0.048	-0.004
ZPERC	0.042	0.043	-0.002
ADIMP	0.035	0.042	-0.007

The surrogate-based estimates consistently lie within the bootstrap confidence intervals of the direct estimates, and the rank ordering of influential parameters is preserved (top five parameters are identical in both analyses). This validation gives confidence that the surrogate-assisted MC Shapley workflow—trained on 4,096 Sobol' samples in approximately 1 minute of regression time—produces results that are statistically indistinguishable from the direct MC Shapley computation on the original model.

Computational note: The direct SAC-SMA evaluation required approximately 50 minutes

(21,000 model evaluations at $N = 1{,}000$, $k_{\max} = 2$). The surrogate completes the same calculation in seconds once trained. For high-dimensional or more expensive models, this cost difference makes surrogate-assisted analysis the only practical option.

6. Computational Cost

The surrogate-based workflow makes the correlated sensitivity analysis tractable for a 14-parameter hydrological model:

- **Training:** 4,096 SAC-SMA evaluations (≈ 60 seconds with Numba JIT)
- **Surrogate regression:** ≈ 58 seconds (streaming OMP-CV with 10-fold CV)
- **MC Shapley (exhaustive, $k_{\max} = 3$):** 1,376 coalitions evaluated with $N = 2{,}000$ samples each (≈ 2.75 million surrogate evaluations, ≈ 20 seconds)
- **Bootstrap:** 500 replications (≈ 45 seconds with Numba acceleration)

Since the surrogate evaluation is effectively free ($\ll 1$ ms), the marginal cost of the correlated analysis beyond the RS-HDMR training is limited to the copula sampling and conditional evaluation—approximately 1–2 minutes total. The surrogate—trained in under a minute—can be reused for *multiple* correlation scenarios (e.g., different catchments, alternative prior assumptions, sensitivity to correlation strength) is a major practical advantage.

7. Conclusions

This case study has demonstrated ShapleyX's capabilities for global sensitivity analysis of a conceptual hydrological model under correlated inputs. The key findings are:

1. **RS-HDMR surrogate accuracy:** A sparse Legendre expansion (338 active terms from 8,876 candidates) achieves $R^2 = 0.989$ on the NSE response surface, trained on only 4,096 Sobol' samples in under a minute of regression time.
2. **Independent-input analysis:** The lower zone free water parameters dominate SAC-SMA's NSE response (LZFPM = 0.249, LZFSM = 0.206, REXP = 0.165). Upper zone parameters are negligible (< 0.02).
3. **Correlated-input analysis:** Using the Vrugt et al. (2006) posterior correlation matrix, the Shapley effects redistribute substantially. PFREE rises from rank 14 to rank 5 ($\Delta = +0.069$), LZPK more than doubles ($\Delta = +0.054$), and LZFPM drops by 33% ($\Delta = -0.081$). The upper zone parameters remain negligible in both analyses.
4. **Methodological insight:** Standard sensitivity analysis assuming independent inputs can produce misleading parameter rankings for models with strong parameter correlations. The Monte Carlo Shapley method with a Gaussian copula provides a principled, practical alternative.
5. **Computational practicality:** The surrogate-based workflow enables correlated sensitivity analysis at negligible marginal cost once the surrogate is trained. Multiple correlation scenarios can be explored without additional model evaluations.

Code Availability

All code and results for this analysis are available in the ShapleyX repository Bennett (2026b)

(`Sac_Example/sac_sma_shapleyx_example.ipynb`).

Data Availability

St Helens Creek forcing and discharge data are provided in the `Sac_Example` directory alongside the notebook (sourced from SILO Jeffrey et al. (2001) and the Bureau of Meteorology Hydrologic Reference Stations Zhang et al. (2014)). The SAC-SMA model implementation is adapted from Markert (2026).

Acknowledgments

The author thanks the National Weather Service Office of Hydrologic Development for the development and public distribution of the SAC-SMA model, and Vrugt et al. for publishing the detailed posterior correlation matrix.

References

S. X. Zhang, M. Bari, G. Amirthanathan, D. Kent, A. MacDonald, and D. Shin. Hydrologic Reference Stations to monitor climate-driven streamflow variability and trends. In *Hydrology and Water Resources Symposium 2014*, page 1048, 2014.

R. J. C. Burnash, R. L. Ferral, and R. A. McGuire. *A Generalized Streamflow Simulation System—Conceptual Modeling for Digital Computers.* US National Weather Service, 1973.

G. Chastaing, F. Gamboa, and C. Prieur. Generalized Hoeffding-Sobol decomposition for dependent variables—application to sensitivity analysis. *Electronic Journal of Statistics*, 6:2420–2448, 2012.

Y. Fan, G. Huang, B. W. Baetz, Y. Li, and K. Huang. Development of a copula-based particle filter (CopPF) approach for hydrologic data assimilation under consideration of parameter interdependence. *Water Resources Research*, 53:4850–4870, 2017.

B. Iooss and C. Prieur. Shapley effects for sensitivity analysis with correlated inputs: comparisons with Sobol' indices, numerical estimation and applications. *International Journal for Uncertainty Quantification*, 9(5), 2019.

A. J. Jakeman and G. M. Hornberger. How much complexity is warranted in a rainfall-runoff model? *Water Resources Research*, 29(8):2637–2649, 1993.

K. Markert. sacsma: Sacramento Soil Moisture Accounting model (Python). `https://github.com/KMarkert/sacsma`, accessed 2026.

V. Koren, M. Smith, and Q. Duan. Use of a priori parameter estimates in the derivation of spatially consistent parameter sets of rainfall-runoff models. In Q. Duan et al., editors, *Calibration of Watershed Models*, volume 6 of *Water Science and Application*, pages 239–254. American Geophysical Union, 2003.

G. Li, S. W. Wang, and H. Rabitz. Practical approaches to construct RS-HDMR component functions. *Journal of Physical Chemistry A*, 106(37):8721–8733, 2002.

A. B. Owen. Sobol' indices and Shapley value. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):245–251, 2014.

A. B. Owen and C. Prieur. On Shapley value for measuring importance of dependent inputs. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):986–1002, 2017.

S. Razavi, M. Mousavi, A. A. Abedini, and M. R. Rafiee. Generalized likelihood uncertainty estimation (GLUE) and SCEM-UA: application to SAC-SMA model. Technical Report MR2007-305, DTU Environment, 2007.

F. Bennett. ShapleyX: A Python package for global sensitivity analysis with RS-HDMR. *Journal of Open Source Software* (submitted), 2026.

F. Bennett. ShapleyX repository. `https://github.com/frbennett/shapleyx`, v0.6.0, 2026.

M. J. Shin, J. H. A. Guillaume, B. F. W. Croke, and A. J. Jakeman. Addressing ten questions about conceptual rainfall-runoff models using a metacognitive modelling approach. *Environmental Modelling & Software*, 39:123–137, 2013.

I. M. Sobol'. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics*, 7(4):86–112, 1967.

I. M. Sobol'. Sensitivity analysis for non-linear mathematical models. *Mathematical Modelling and Computational Experiment*, 1:407–414, 1993.

E. Song, B. L. Nelson, and J. Staum. Shapley effects for global sensitivity analysis: Theory and computation. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):1060–1083, 2016.

J. A. Vrugt, H. V. Gupta, W. Bouten, and S. Sorooshian. A shuffled complex evolution metropolis algorithm for optimization and uncertainty assessment of hydrologic model parameters. *Water Resources Research*, 39(8):1201, 2003.

J. A. Vrugt, H. V. Gupta, S. C. Dekker, S. Sorooshian, T. Wagener, and W. Bouten. Application of stochastic parameter optimization to the Sacramento Soil Moisture Accounting model. *Journal of Hydrology*, 325:288–307, 2006.

Queensland Department of the Environment, Tourism, Science and Innovation. WetlandInfo: St Helens Creek (Mackay Whitsunday) Water Quality Improvement Plan catchment—facts and maps. `https://wetlandinfo.detsi.qld.gov.au/wetlands/facts-maps/wqip-st-helens-creek-mackay-whitsunday/`, accessed 2025.

S. J. Jeffrey, J. O. Carter, K. B. Moodie, and A. R. Beswick. Using spatial interpolation to construct a comprehensive archive of Australian climate data. *Environmental Modelling & Software*, 16(4):309–330, 2001.

Z. Zhang, V. Koren, S. Reed, M. Smith, and B. Cosgrove. SAC-SMA a priori parameter differences and their impact on distributed hydrologic model simulations. *Journal of Hydrology*, 408(3–4):199–212, 2011.

Affiliation:

Frederick R. Bennett
Department of Environment, Tourism, Science and Innovation
Queensland Government
GPO Box 2454
Brisbane QLD 4001, Australia
E-mail: frederick.bennett@detsi.qld.gov.au