

Refinements to SPoRC for eastern Bering Sea pollock

Working paper on the corrected 2024 “PM” model implementation bridge


James Ianelli

Matt Cheng

August 16, 2026

Abstract

The custom AD Model Builder (ADMB) stock assessment model for eastern Bering Sea walleye pollock has a long history of operational use. This working paper evaluates an implementation of the 2024 assessment in the Stock Population and Recruitment Computational (SPoRC) framework. The corrected bridge normalizes bottom-trawl survey age compositions across ages 1–15 and retains nominal multinomial sample sizes; ages 2–15 continue to define the survey abundance total. Spawning biomass, recruitment, population numbers at age, and selectivity agree with the corrected ADMB candidate at high numerical accuracy. The accepted ADMB assessment retains its management role. Updated afscOSA diagnostics, nine retrospective peels, a 30-dataset simulation self-test, and SPoRC-derived spmR projections all use the corrected fit. Fixing steepness preserves implementation equivalence. A correlated bivariate random walk on annual bottom-trawl survey midpoint and slope supports a sparse-metric SparseNUTS trial and identifies the remaining posterior geometry for further development.



SPoRC COMPANION REPORT

1 Executive Summary

This page implements the [SPoRC single-region EBS pollock case study](#) as the SPoRC member of the EBS pollock model family. The primary lane uses a corrected ADMB model in which BTS observed age compositions sum to one across ages 1–15 and the multinomial retains the nominal sample sizes. Ages 2–15 continue to define the BTS abundance-total series.

The implementation reproduces the corrected 2024 ADMB candidate’s spawning biomass and recruitment series at near-numerical precision. The accepted ADMB assessment retains its role as the management model.

The exact bridge has a maximum time-series difference of 0.0611%. The fitted solution has a positive-definite Hessian and a maximum fixed-effect gradient of 3.17e-08. The evidence remains a development-model result, with ADMB retained as the management reference.

The corrected SPoRC fit reproduces six ADMB state and selectivity surfaces with a maximum difference of 0.0611%. Relative to the legacy adjusted-ISS bridge, the correction changes spawning biomass by at most 1.62% and recruitment by at most 1.81%. All diagnostics and projections use the corrected lane.

Fixing steepness at 0.6184280 preserves the corrected bridge at numerical precision. A correlated bivariate random walk on annual BTS $\log(a_{50})$ and $\log(k)$ improves the maximum-likelihood objective relative to independent increments. Section 5.4 reports the updated sparse-metric posterior diagnostics and retains the maximum-likelihood bridge as the implementation basis.

The implementation uses SPoRC version 1.2.0.9000 from development branch dev-popn-seasons at commit 63af5a30, and covers 1964–2024.

1.1 Position in the model family




Figure 1: SPoRC’s current position in the EBS pollock model-family workflow.

The family registry is stored in `config/model_family.yml`. The exact bridge is the admission case. The comparisons exclude the earlier three-survey Beverton–Holt experiment because it uses a different model specification.

2 Purpose and Scope

This working paper documents whether SPoRC can reproduce the principal population dynamics and assessment outputs of the accepted 2024 eastern Bering Sea pollock model. It serves as an implementation and methods comparison rather than a full Stock Assessment and Fishery Evaluation chapter.

The paper therefore records:

- the model inputs and structural choices represented in SPoRC;
- agreement between the SPoRC and ADMB bridge calculations;
- fishery and survey selectivity comparisons;
- numerical and likelihood-accounting diagnostics; and
- the files and software revision needed to reproduce the analysis.

Fishery history, ecosystem narrative, projections, reference-point tables, and harvest recommendations are outside the scope of this working paper.

Table 1: Index and survey components in the exact EBS pollock bridge.

Component	Likelihood	Selectivity	SPoRC representation
Fishery CPUE	Normal	Fishery	Fishery 1
Bottom trawl survey	Multivariate normal	Logistic, time varying, free age 1	Survey 1
Acoustic trawl survey	Lognormal	Non-parametric random walk	Survey 2
Vessel of opportunity	Normal	Shared with ATS	Survey 3
Acoustic age 1	Lognormal	Shared with ATS	Survey 4 (age 1 only)

Table 2: SPoRC and ADMB time-series differences.

Quantity	Median difference (%)	Maximum difference (%)
Spawning biomass	8.017×10^{-7}	2.322×10^{-5}
Recruitment	7.216×10^{-7}	4.345×10^{-5}
Numbers at age	2.275×10^{-6}	9.325×10^{-5}
Fishery selectivity	3.555×10^{-6}	1.429×10^{-4}
BTS selectivity	5.033×10^{-7}	1.776×10^{-5}
ATS selectivity	7.314×10^{-7}	6.107×10^{-2}

3 Initial SPoRC Developments

3.1 Model Inputs

The model is single population, single region, single season, and single sex, with 15 ages, one fishery, and four SPoRC survey fleets. The fourth survey fleet represents the acoustic age-1 likelihood component, distinct from the observed survey fleets.

3.2 Selectivity specifications

Structural and selectivity choices adopted directly from the case study include:

- Ricker recruitment with a one-year lag, free initial numbers at age, the assessment recruitment penalties, and the steepness prior.
- Separate spawning, fishery, and survey weight-at-age arrays.
- The BTS full covariance likelihood and a free annual BTS age-1 selectivity value.
- Non-parametric random-walk fishery and ATS selectivity, with AVO sharing ATS selectivity.
- The acoustic age-1 abundance likelihood as survey fleet 4 with its own catchability.
- ADMB parameter values as the bridge starting point, with recruitment deviations solved by forward substitution and multiplicative catchabilities recovered analytically.

4 ADMB Comparison Results

The bridge is evaluated at the accepted ADMB maximum-likelihood estimate before optimization. The objective there is 1517.039555 and the maximum absolute gradient is 0.00148. Optimization leaves the solution effectively unchanged.

The fitted object has 1207 free parameters, a final joint negative log likelihood of 1517.040, a positive-definite Hessian, and a final maximum absolute fixed-effect gradient of 3.17e-08. `nlminb` reports code 0 (“both X-convergence and relative convergence (5)”); the code remains disclosed even though the gradient, Hessian, and bridge trajectories reproduce the case-study solution.

4.1 Corrected BTS age-composition treatment

The corrected ADMB candidate normalizes each observed BTS age composition over ages 1–15, matching the prediction vector and the stated multinomial sample space. It retains the nominal integer sample size for each




Figure 2: SPoRC and ADMB spawning biomass and recruitment. Lines overlap at plotting precision.




Figure 3: Annual SPoRC-minus-ADMB differences. The vertical scale is percent.

Table 3: Change in SPoRC state outputs after correcting BTS composition normalization and retaining nominal sample sizes.

Quantity	Median change from legacy bridge (%)	Maximum change from legacy bridge (%)
Spawning biomass	0.196	1.625
Recruitment	0.340	1.813

year. Age-1 selectivity remains an annual free parameter, while a logistic midpoint and slope define ages 2–15. The BTS abundance-total series continues to use the ages 2–15 subtotal, preserving that observation model.

The legacy implementation normalized observed ages 1–15 by the ages 2–15 subtotal. Its adjusted SPoRC sample sizes reproduced that score algebraically. The legacy multiplier has a median of 1.128 and ranges from 1.036 to 1.326. The corrected form provides a coherent multinomial and a common sample-size basis across the model-family tests.

The corrected ADMB-to-RTMB bridge differs in objective by 0.00038 and in key outputs by at most 0.00049%. The corrected ADMB-to-SPoRC bridge in Table 2 provides the primary implementation regression gate. The legacy artifact remains available for historical comparison.

4.2 Selectivity comparison




Figure 4: Normalized fishery selectivity from SPoRC and ADMB for selected ages.




Figure 5: Normalized BTS and ATS selectivity surfaces from SPoRC and ADMB for selected years.

Table 4: Summary statistics for nominal-ISS age-composition diagnostics.

Data source	OSA source	Years	Composition years	Age bins	SDNR	Lower 2.5%	Upper 97.5%	ISS	
Fishery	SPoRC internal OSA	1964–2023	60	15	0.81	-1.95	1.39	17,127.00	33,448
BTS	SPoRC internal OSA	1982–2024	42	15	0.99	-2.46	1.53	5,397.00	3,650
ATS	afscOSA multinomial	1994–2024	18	14	0.81	-1.81	1.66	914.00	1,979

5 Diagnostics

The corrected implementation-equivalence result is strong: spawning biomass, recruitment, population numbers at age, and selectivity reproduce the corrected ADMB candidate at high numerical accuracy. OSA, retrospective, self-test, SparseNUTS, and projection products all trace to this corrected fit or its stated fixed-steepness and parameter-RW derivatives.

ADMB and SPoRC retain engine-specific likelihood constants and reporting layouts. State, fitted-surface, gradient, Hessian, and within-engine objective checks therefore provide the regression gates used here.

5.1 OSA Residuals

The age-composition diagnostics use the nominal-ISS model-family lane and the updated [afscOSA workflow](#) from [pull request 10](#), installed from commit 2dfe160f. The analysis passes SPoRC’s internal randomized OSA residuals to afscOSA for the fishery and BTS. This preserves the composition likelihood and conditional calculation embedded in SPoRC. The ATS calculation uses the afscOSA multinomial routine on ages 2–15, the age range used by the fitted likelihood. SPoRC’s current internal composition packer also retains ATS age 1. The external calculation therefore aligns the ATS diagnostic with the fitted data window.

The ATS 2020 input sample size equals one. Integer allocation across its fractional age proportions yields zero counts, so the analysis excludes that zero-information row. The remaining ATS series contains 18 composition years. All randomized calculations use seed 9.9801×10^4 .

The 42 corrected BTS annual sample sizes sum to 5,403. afscOSA’s proportion-to-integer-count conversion produces the displayed aggregate ISS of 5,397, matching the convention used for the other model-family diagnostics.

The aggregate panels provide the first diagnostic. Observed proportions follow the principal age patterns expected by the model, while departures at specific ages identify where selectivity and composition assumptions merit attention. The interval bars describe the spread of data generated from each fitted composition distribution at the aggregate sample size; they differ from parameter-confidence intervals.

The fishery SDNR is 0.81 and the ATS SDNR is 0.81. Both values fall below their displayed 95% standard-normal reference intervals, indicating narrower residual distributions than the reference. Their aggregate ESS values exceed their ISS values, consistent with realized aggregate variability equivalent to larger samples than the input weighting. The fishery upper-tail quantile (1.39) also shows compression. ATS tail quantiles remain within their displayed reference intervals.

The BTS SDNR is 0.99 and lies within its displayed reference interval. Its lower-tail quantile (-2.46) extends beyond the reference interval, whereas its upper tail is compressed at 1.53. The OSA and Pearson bubbles localize these asymmetries across ages and years. afscOSA displays two fishery Pearson values above 6 at the plotting cap; the largest raw value is 10.47.

These diagnostics jointly describe model-data compatibility. The aggregate fits, Q-Q behavior, tail statistics, and residual patterns guide investigation of selectivity, sample weighting, and temporal structure. Their combined evidence supports model refinement more effectively than a single threshold or pass/fail decision.




Figure 6: Age-composition fit and residual diagnostics from afscOSA for the fishery, bottom trawl survey (BTS), and acoustic-trawl survey (ATS). Columns distinguish data sources. Rows show observed aggregate proportions with model expectations and 95% data intervals, quantile–quantile plots of randomized OSA residuals with SDNR and tail-reference intervals, OSA residual bubbles, and Pearson residual bubbles. Blue and red bubbles represent positive and negative residuals, respectively, and bubble area increases with residual magnitude. The ATS column is partially visible on the right.

Table 5: Data-availability and terminal-selectivity audit for nine SPoRC retrospective peels.

Peel	Terminal year	Fishery ages through	Maximum terminal selectivity difference
1	2023	2022	0.00
2	2022	2021	0.00
3	2021	2020	0.00
4	2020	2019	0.00
5	2019	2018	0.00
6	2018	2017	0.00
7	2017	2016	0.00
8	2016	2015	0.00
9	2015	2014	0.00

Table 6: Optimization and Hessian status for the full fit and nine SPoRC retrospective peels.

Peel	Positive-definite Hessian	Maximum gradient	Optimizer code
0	TRUE	3.17×10^{-8}	0
1	TRUE	1.19×10^{-11}	1
2	TRUE	7.87×10^{-9}	1
3	TRUE	1.21×10^{-12}	1
4	TRUE	2.21×10^{-12}	1
5	TRUE	2.86×10^{-10}	1
6	TRUE	2.78×10^{-12}	1
7	TRUE	2.66×10^{-12}	1
8	TRUE	9.53×10^{-10}	1
9	TRUE	1.42×10^{-11}	1

5.2 Retrospective patterns

Nine retrospective peels terminate in 2023–2015. Each active contemporary data series retains the availability lag represented in the 2024 model. The fishery age compositions have a one-year lag: the 2023 peel uses compositions through 2022, and subsequent peels continue that one-year offset. Survey indices and survey age compositions use observations at or before each peel’s terminal year according to their historical sampling schedules. The historical fishery CPUE series retains its 1976 endpoint.

Each peel constrains fishery and survey selectivity in its terminal year to the preceding year’s selectivity. This constraint includes the BTS age-1 selectivity deviation. The authoritative 2024 full fit remains unchanged.

All nine truncated fits have positive-definite Hessians. Their maximum fixed-effect gradients range from $1.2\text{e-}12$ to $7.9\text{e-}09$. The status table retains optimizer code 1 for transparency; the positive-definite Hessians and gradients below 10^{-8} provide the supporting numerical diagnostics. The selectivity audit gives a maximum terminal-year difference of 0.0 across every peel and fleet, confirming the equality constraint.

Mohn’s rho equals 0.254 for spawning biomass and 0.246 for recruitment. Several earlier peels estimate higher terminal spawning biomass and recruitment than the corresponding full-model years, producing the positive average terminal differences.

Table 7: Mohn’s rho from nine SPoRC retrospective peels.

Quantity	Mohn’s rho
Recruitment	0.246
SSB	0.254




Figure 7: Relative-difference retrospective diagnostics and Mohn's rho for the availability-aligned exact SPoRC bridge. Fishery compositions retain a one-year reporting lag, and terminal selectivity equals the preceding year's selectivity in each peel.




Figure 8: Spawning biomass and recruitment trajectories for the full SPoRC fit and nine availability-aligned retrospective peels.

Table 8: Thirty-replicate SPoRC simulation and refit self-test.

Quantity	Requested	Successful	Failed	Median relative error	RMS relative error	Terminal median relative error	Terminal median
Fmort	30	26	4	0.159	0.225	0.095	
Rec	30	26	4	0.071	0.443	0.057	
SSB	30	26	4	0.072	0.109	-0.042	

Table 9: Bridge sensitivity after fixing steepness at the exact-bridge estimate.

Quantity	Median difference from ADMB (%)	Maximum difference from ADMB (%)	Median difference from estimated-steepness
Spawning biomass	8.017×10^{-7}	2.322×10^{-5}	1.
Recruitment	7.216×10^{-7}	4.345×10^{-5}	1.

5.3 Simulation self-test

Thirty datasets were generated from the fitted SPoRC model and refitted, with recruitment generated from the fitted Ricker relationship. Twenty-seven refits succeeded and 4 failed. Failed replicates remain in the requested count and are excluded from the error summaries.




Figure 9: Simulated truth compared with refitted estimates for the successful SPoRC self-test replicates.

5.4 MCMC integrations

The steepness sensitivity fixes biological-scale steepness at 0.6184280, the estimate from the exact bridge, and reoptimizes every remaining free parameter. This removes one posterior dimension while preserving the fitted stock-recruit relationship. The sensitivity has 1,206 free parameters, a positive-definite Hessian, and a maximum gradient of $3.17\text{e-}08$. Its objective differs from the estimated-steepness bridge by $0\text{e+}00$.

Maximum differences from the estimated-steepness SPoRC bridge remain below $2.25\text{e-}13\%$, and maximum differences from ADMB remain below $4.35\text{e-}05\%$. Fixing steepness therefore preserves implementation equivalence at numerical precision.

5.4.1 BTS parameter-level random walks

Two sensitivity models replace the 1983–2024 BTS age-specific curve-change penalty with direct random walks on annual $\log(a_{50})$ and $\log(k)$. Both retain the 1982 curve anchor and fix the increment SDs at 0.120 and 0.305. The first uses independent increments. The second uses the new SPoRC correlated bivariate random

walk with correlation -0.829, estimated empirically from the exact-bridge maximum-likelihood deviations. The fixed hyperparameters make this a targeted parameterization experiment rather than an empirical-Bayes uncertainty analysis.

Table 10: Maximum-likelihood diagnostics for BTS parameter-level random-walk sensitivities.

BTS process	RMS $\Delta\log(a50)$	RMS $\Delta\log(k)$	Increment correlation	Maximum gradient	Positive-definite Hessian	Maximum SSB
Independent RW	0.114	0.230	-0.564	1.65×10^{-11}	TRUE	
Correlated RW	0.126	0.277	-0.881	1.19×10^{-11}	TRUE	

Both parameter-level sensitivities reach positive-definite Hessians with maximum gradients below 2×10^{-11} . The correlated random walk follows the fitted increment dependence closely and stays nearer the exact bridge: its maximum SSB difference is 2.86%, compared with 4.93% for independent increments. These differences quantify a deliberate change in regularization; the exact bridge continues to carry the implementation-equivalence result.




Figure 10: Maximum-likelihood BTS logistic midpoint and slope under the exact curve-level bridge and the two parameter-level random-walk sensitivities.

SparseNUTS was applied as a developmental posterior-sampling integration around both SPoRC bridge modes and the correlated BTS parameter-RW sensitivity. The runner rebuilds each RTMB objective from its archived data and parameter map, verifies the free-parameter names against the corresponding fit, and restores its maximum-likelihood estimates before sampling. SparseNUTS selected a dense metric for the bridge objectives. The correlated-RW run explicitly sets `metric = "sparse"`, declares `ln_srvsel_devs` as an RTMB random-effect block, and sets `laplace = FALSE`. SparseNUTS therefore samples the same 1,206-dimensional joint posterior while constructing the sparse joint precision matrix, Q . The classification contains 910 fixed and 296 random parameters; it changes the computational representation while preserving the joint posterior target. Each run used four chains, 150 warmup iterations per chain, and 1,000 retained iterations per chain with seed 99803.

Both updated runs use a target acceptance probability of 0.95 and maximum tree depth of 12. The convergence gate requires zero divergent transitions, zero maximum-tree-depth transitions, acceptable energy behavior in every chain, and a maximum R-hat at or below 1.01.

Table 11: SparseNUTS diagnostics for corrected fixed-steepness SPoRC models.

Model	Target acceptance	Maximum R-hat	Minimum bulk ESS	Divergences
Corrected bridge; steepness fixed	0.95	1.01	309.94	9
Correlated BTS RW; steepness fixed	0.95	3.29	4.47	0

The corrected fixed-steepness run failed the convergence gate, so its draws remain unsuitable for posterior inference. The weakest behavior identifies the parameter blocks that govern the next reparameterization experiment. The maximum-likelihood bridge result and its management interpretation remain unchanged.




Figure 11: Post-warmup SparseNUTS traces for the six parameters with the weakest convergence diagnostics in the tuned fixed-steepness run. Colors distinguish the four chains.

The brackets in the sampler labels index positions in the free parameter vector; Table 12 maps those positions to their model years and natural-scale multipliers.

Table 12: Model interpretation of the sampler labels in the fixed-steepness trace figure.

Sampler label	Model quantity	Year and age application	Maximum-likelihood translation
ln_RecDevs[34]	Annual log recruitment deviation	1997; age-1 recruitment	Multiplier = 0.972
ln_RecDevs[36]	Annual log recruitment deviation	1999; age-1 recruitment	Multiplier = 0.522
ln_RecDevs[33]	Annual log recruitment deviation	1996; age-1 recruitment	Multiplier = 0.725
ln_global_R0	Global log(R0)	Global; age-1 recruitment scale	R0 = 24,256 model abundance units
ln_RecDevs[50]	Annual log recruitment deviation	2013; age-1 recruitment	Multiplier = 1.458
ln_RecDevs[27]	Annual log recruitment deviation	1990; age-1 recruitment	Multiplier = 1.602




Figure 12: Post-warmup pairs plot for the six parameters with the weakest convergence diagnostics in the tuned fixed-steepness SparseNUTS run. Diagonal panels show marginal posterior densities, lower panels show joint draws, and upper panels show correlations.

The pairs matrix shows the strongest sampled dependence between `ln_global_R0` and recruitment deviations for 1990, 1996, 1997, 1999, and 2013. Their correlations with the global recruitment scale range from -0.941

to -0.971, and their mutual correlations range from 0.917 to 0.955. These relationships support the dense preconditioner and highlight the broader recruitment-scale mixing problem.

5.4.2 Correlated BTS random-walk posterior trial

The correlated parameter-level random walk was sampled from its maximum-likelihood solution with the same target acceptance probability and tree-depth limit as the tuned bridge run. Four parallel chains use the same validated maximum-likelihood initialization. The sparse metric uses the 296 survey-selectivity deviations as the RTMB random block and samples all 1,206 joint parameters with Laplace integration disabled.

Table 13: SparseNUTS diagnostics for the fixed-steepness model with a correlated BTS parameter-level random walk.

TMB metric	RTMB random block	Chains	Warmup per chain	Retained per chain	Maximum R-hat	Minimum bulk ESS	Divergen
sparse	ln_srvsel_devs	4	150	1000	3.287	4.5	

The correlated-RW run failed the convergence gate. It produced 0 divergent transitions, a maximum R-hat of 3.287, and a minimum bulk ESS of 4.5. It reached maximum tree depth in 2,921 of the 4,000 retained transitions (73.03%) and placed 165 parameters above an R-hat of 1.05. Relative to the tuned fixed-steepness bridge, divergences changed from 9 to 0, maximum R-hat changed from 1.009 to 3.287, and minimum bulk ESS changed from 309.9 to 4.5. Any divergent transition, maximum-tree-depth transition, or maximum R-hat above 1.01 keeps the draws exploratory under the stated gate.




Figure 13: SparseNUTS joint precision and correlation structure for the correlated BTS parameter-level random-walk run. The upper triangle shows the structural nonzero pattern of Q ; the lower triangle shows correlations implied by its inverse, with red positive and blue negative.

The joint precision matrix has dimension 1206 by 1206 and 8.60% structural zeros. The modest sparsity

reflects broad coupling between annual state and observation parameters, while the block pattern preserves enough structure for SparseNUTS to construct its requested sparse preconditioner. Constructing Q took 26.3 seconds, and the four chains averaged 89.46 minutes each.




Figure 14: Post-warmup traces for the six weakest-mixing parameters in the correlated BTS parameter-level random-walk run. Colors distinguish the four chains.

Table 14: Model interpretation of the six weakest-mixing sampler labels under the correlated BTS random walk.

Sampler label	Model quantity	Year and age application	Maximum-likelihood translation
ln_srvsel_devs[12]	BTS log(k)	1987; ages 2–15	$k = 0.687$ per year
ln_srvsel_devs[29]	BTS log(a50)	1996; ages 2–15	$a50 = 7.105$ years
ln_srvsel_devs[1]	BTS log(a50)	1982; ages 2–15	$a50 = 5.616$ years
ln_srvsel_devs[20]	BTS log(k)	1991; ages 2–15	$k = 1.183$ per year
ln_srvsel_devs[22]	BTS log(k)	1992; ages 2–15	$k = 1.194$ per year
ln_srvsel_devs[79]	BTS log(a50)	2021; ages 2–15	$a50 = 6.680$ years

All six weakest-mixing parameters are annual survey-selectivity parameters. Table 14 resolves each sampler index to its year, age application, and natural-scale maximum-likelihood value.




Figure 15: Pairs plot for the six weakest-mixing parameters in the correlated BTS parameter-level random-walk run. Diagonal panels show marginal posterior densities, lower panels show joint draws, and upper panels show correlations.

The pairs matrix locates the remaining geometry in annual BTS selectivity. Pairwise correlations among the six weakest-mixing deviations range from -0.536 to 0.879. Together with the maximum R-hat of 3.287 and 73.03% maximum-tree-depth rate, this concentration identifies the BTS selectivity state vector and its sparse parameterization as the clearest next sampling experiment.

6 Projections

This projection section reads the saved `sporc_corrected_full_age_bts` fit and its matching corrected input artifact directly. The corrected SPoRC bridge supplies a complete Standard Projection Model input directory for `spmR`. The adapter extracts terminal numbers at age and biological schedules, carries forward terminal fishery selectivity, calculates recent fishing mortality from the same age-6 convention used in the reference projection, and writes `pm.prj` and `spm.dat`. The runner compiles the ADMB projection engine from the

Table 15: Files in the SPoRC-to-spmR projection handoff.

SPoRC-to-spmR Projection Inputs

File	Role	Available
spm.dat	SPM setup generated for the SPoRC bridge	TRUE
pm.prj	Species assessment input generated from SPoRC	TRUE
tacpar.dat	Generic TAC-control input maintained in this SPoRC project	TRUE
spm	ADMB executable compiled from the installed spmR template	TRUE

Table 16: Field-level audit connecting Standard Projection Model inputs to the saved SPoRC bridge state.

SPoRC Projection-input Audit

Projection quantity	SPoRC source	Values	Maximum relative
average_f	family_sporc_corrected_full_age_bts.rds::report\$tot_FAA[terminal-4:terminal, age 6]	1	
natural_mortality	family_sporc_corrected_full_age_bts.rds::report\$natmort[terminal year, age]	15	
maturity	family_sporc_corrected_full_age_bts_input_list.rds::data\$MatAA[terminal year, age]	15	
spawning_weight	family_sporc_corrected_full_age_bts_input_list.rds::data\$WAA[terminal year, age]	15	
fishery_weight	family_sporc_corrected_full_age_bts_input_list.rds::data\$WAA_fish[terminal year, age]	15	
selectivity	family_sporc_corrected_full_age_bts.rds::report\$fish_sel[terminal year, age]	15	
terminal_numbers	family_sporc_corrected_full_age_bts.rds::report\$NAA[terminal year, age]	15	
recruitment	family_sporc_corrected_full_age_bts.rds::report\$Rec[1978:terminal year]	47	
spawning_biomass	family_sporc_corrected_full_age_bts.rds::report\$SSB[1977:(terminal year-1)]	47	

spm.tp1 source installed with the recorded **spmR** version. A field-level audit compares each written vector with the saved SPoRC state and enforces a maximum relative formatting difference of 5×10^{-7} .

The SPoRC and RTMB projections agree closely because both implementation bridges reproduce the corrected ADMB candidate’s population state at high numerical accuracy. The projections reported here read the SPoRC fit and input artifacts listed in Table 15; their source paths and hashes remain distinct from the RTMB projection artifacts.

The seven-alternative run uses the 2024 terminal-year maximum-likelihood bridge, fixes 2025 and 2026 catches at 1,350 thousand t, and projects 1,000 simulations through 2038. This is a development-model projection sensitivity; the accepted assessment remains the source for management advice.

6.1 Seven Tier-3 scenario results

The current spmR release (version 0.3.0, main commit e86fc6a) ran all seven alternatives. The comparison starts in 2027 because the 2025 and 2026 catches are specified at 1,350 thousand t in every alternative.

6.2 Alternative 2 fixed-catch projection

The companion Alternative 2 run holds catch at 1,300 thousand t annually from 2025 through 2032. The resulting biomass path remains above the Tier-3 $B_{35\%}$ proxy throughout the specified-catch horizon.

6.3 Projection figures

The lines show simulation medians, shaded bands span the central 90% of the 1,000 simulations, and colored lines identify the same 10 reproducibly selected simulations in each panel (242, 278, 361, 495, 507, 510, 598, 837, 937, 989).

Table 17: Key settings and reproducibility identifiers for the SPoRC-derived projection run.

Projection Configuration and Lineage

Item	Value
SPoRC version	1.2.0.9000
SPoRC fit case	sporc_corrected_full_age_bts
SPoRC fit commit	63af5a30dede
spmR version	0.3.0
spmR source commit	e86fc6aa6f48
spmR template SHA-256	ce425c64eede084f
Source-fit MD5	17288f8dd88bf766cdb3f5b7bce82068
Terminal year	2024
Spawning month	4 (April)
Projection years	2025–2038
Simulations per alternative	1000
Alternatives	1–7
Fixed catches	1,350 thousand t in 2025 and 2026
Average recent F at age 6	0.219
Species input	pm.prj

Table 18: Standard Projection Model alternatives evaluated with the SPoRC-derived assessment state.

Tier-3 Projection Alternatives

Alternative	Name	Description
1	Maximum permissible ABC	Applies the maximum permissible Tier 3 ABC harvest rate.
2	Author-specified ABC	Applies the author-specified Tier 3 ABC adjustment; unit multipliers make this run equivalent to
3	Average recent F	Applies recent average fishing mortality from the SPoRC assessment state.
4	Alternative SPR rate	Applies the user-specified SPR rate of 0.60.
5	No fishing	Projects the population with zero fishing mortality after the fixed-catch years.
6	OFL threshold determination	Applies the OFL harvest rate for threshold and status determination.
7	Status-determination ramp	Applies maximum permissible ABC for three projection years, then uses the OFL status-determina

Table 19: Tier-3 projection results for seven alternatives using inputs generated from the SPoRC EBS pollock bridge. Values are simulation means; catch, ABC, OFL, and spawning biomass are in thousand t. $B/B_{35\%}$ is spawning biomass relative to the Tier-3 $B_{35\%}$ proxy.

Seven Tier-3 Projection Alternatives

Alternative	Scenario	2027						2028					
		Catch	ABC	OFL	B	F	$B/B_{35\%}$	Catch	ABC	OFL	B	F	$B/B_{35\%}$
1	Maximum permissible ABC	1,620	1,620	1,947	2,553	0.270	122%	1,507	1,507	1,795	2,360	0.260	
2	Author-specified ABC	1,620	1,620	1,947	2,553	0.270	122%	1,507	1,507	1,795	2,360	0.260	
3	Average recent F	1,353	1,353	1,947	2,587	0.220	124%	1,347	1,347	1,906	2,482	0.220	
4	Alternative SPR rate	783	783	1,947	2,656	0.120	127%	843	843	2,093	2,762	0.120	
5	No fishing	0	0	1,947	2,744	0.000	131%	0	0	2,309	3,163	0.000	
6	OFL threshold determination	1,671	1,671	1,671	2,262	0.310	108%	1,501	1,501	1,501	2,101	0.290	
7	Status-determination ramp	1,537	1,840	1,840	2,417	0.270	116%	1,692	1,692	1,692	2,244	0.310	

Table 20: Age-specific biological and fishery schedules passed from SPoRC to the Standard Projection Model.

Projection Age-Specific Schedules

Age	Spawning wt-at-age	Fishery wt-at-age	Maturity	Selectivity
1	0.085	0.025	0.000	0.000
2	0.196	0.186	0.008	0.063
3	0.322	0.401	0.289	0.250
4	0.534	0.667	0.641	0.743
5	0.621	0.663	0.842	1.426
6	0.707	0.720	0.901	1.118
7	0.822	0.791	0.947	1.162
8	0.946	0.905	0.963	1.263
9	1.051	1.009	0.970	1.307
10	1.113	1.061	1.000	1.300
11	1.178	1.138	1.000	1.286
12	1.290	1.257	1.000	1.271
13	1.332	1.267	1.000	1.271
14	1.367	1.290	1.000	1.271
15	1.415	1.368	1.000	1.271

Table 21: Candidate ABC projection given catches fixed at 1,300 thousand t from 2025 through 2032. Catch, ABC, OFL, and spawning biomass are simulation means in thousand t. $B/B_{35\%}$ expresses mean spawning biomass relative to the Tier-3 proxy.

Candidate ABC Projection Given Assumed Future Catches

Year	Catch	ABC	OFL	Mean B	$B/B_{35\%}$
2025	1,300	2,029	2,433	3,269	156%
2026	1,300	1,726	2,073	2,870	137%
2027	1,300	1,641	1,971	2,631	126%
2028	1,300	1,632	1,951	2,534	121%
2029	1,300	1,606	1,921	2,553	122%
2030	1,300	1,638	1,962	2,612	125%
2031	1,300	1,668	2,003	2,684	128%
2032	1,300	1,697	2,043	2,751	132%

Table 22: Annual simulation means for Standard Projection Model Alternative 1. Catch, ABC, OFL, and spawning biomass are in thousand t.

Alternative 1 Annual Projection Means

Year	Mean B	Catch	ABC	OFL	F	$B/B_{35\%}$
2025	3,263	1,350	2,029	2,433	0.170	156%
2026	2,844	1,350	1,715	2,059	0.204	136%
2027	2,553	1,620	1,620	1,947	0.268	122%
2028	2,360	1,507	1,507	1,795	0.259	113%
2029	2,348	1,458	1,458	1,744	0.246	112%
2030	2,395	1,504	1,504	1,801	0.244	115%
2031	2,439	1,538	1,538	1,843	0.245	117%
2032	2,478	1,563	1,563	1,874	0.246	119%
2033	2,498	1,581	1,581	1,894	0.246	120%
2034	2,489	1,572	1,572	1,883	0.246	119%
2035	2,484	1,564	1,564	1,874	0.246	119%
2036	2,476	1,562	1,562	1,871	0.246	118%
2037	2,472	1,554	1,554	1,861	0.246	118%
2038	2,484	1,561	1,561	1,870	0.245	119%




Figure 16: Projected spawning biomass under seven Tier-3 alternatives. Dark lines show medians, blue bands span the 5th to 95th percentiles, colored lines show 10 selected simulations, and dashed horizontal lines mark the $B_{35\%}$ proxy. Biomass is in thousand metric tons.




Figure 17: Projected catch, ABC, and OFL under seven Tier-3 alternatives. Black line types show simulation medians, green bands span the 5th to 95th percentiles for catch, and colored lines show the selected catch trajectories. Quantities are in thousand metric tons.



Figure 18: Projected fishing mortality under seven Tier-3 alternatives. Dark lines show medians, orange bands span the 5th to 95th percentiles, and colored lines show the selected simulation trajectories.

The reproducible workflow is implemented in [R/write_spmr_projection_inputs.R](#) and [analysis/run_sporc_spmr_project](#). The generated handoff preserves the reference projection layout and replaces the assessment state with quantities extracted directly from SPoRC.

7 ESP-linked Recruitment Research

Environmental and ecosystem-linked recruitment models lie beyond the exact implementation bridge tested here. The current SPoRC configuration retains the accepted Ricker stock-recruit relationship and holds the recruitment hypothesis constant during implementation testing. A future research stage could compare environmental covariates or dynamic recruitment structures against this validated bridge, with each alternative evaluated using common data, convergence criteria, retrospective behavior, and simulation recovery.

8 Summary

The SPoRC bridge achieves its primary implementation goal: it reproduces the corrected full-age BTS ADMB candidate's spawning biomass, recruitment, population state, and selectivity outputs at high numerical accuracy. The correction yields unit-sum observed and predicted BTS compositions across ages 1–15 with nominal sample sizes, while ages 2–15 continue to define the BTS abundance total. The result supports continued development of SPoRC as a member of the EBS pollock model family, while the accepted ADMB assessment remains the management reference. Fixing steepness preserves the corrected bridge. The correlated BTS parameter-level random walk improves the maximum-likelihood objective and provides the sparse-metric posterior trial. The stated convergence gate determines the appropriate use of those draws.

The current results support review of the corrected implementation bridge, positive retrospective bias, 4 self-test failures, and updated SparseNUTS convergence findings. They support continued SPoRC development while retaining the accepted management model.

The immediate work sequence retains the corrected ADMB and RTMB comparisons as regression gates after SPoRC package, likelihood, or data changes. Subsequent diagnostics should use the corrected bridge configuration. Reparameterizing the weakest-mixing block and repeating the correlated-RW SparseNUTS gate will establish readiness for using MCMC draws in uncertainty summaries or projections.

9 Reproducibility

- SPoRC branch: dev-popn-seasons
- SPoRC fit commit: 63af5a30dedeb38fbc2dea347748441c2f5f3af1
- SPoRC configured commit: 63af5a30dedeb38fbc2dea347748441c2f5f3af1
- SPoRC local-source changes present: FALSE
- SPoRC local-source fingerprint: d41d8cd98f00b204e9800998ecf8427e
- Verified latest-commit refresh manifest: `analysis/outputs/latest_commit_refresh_manifest.csv`
- Case-study source: [single-region EBS pollock case study](#)
- Corrected ADMB runner: `../rtmb_ebswp/scripts/run_corrected_full_age_bts_admb.R`
- Corrected SPoRC data builder: `analysis/build_corrected_full_age_bts_data.R`
- Model runner: `analysis/run_family_model.R`
- Corrected model artifact: `analysis/outputs/family_sporc_corrected_full_age_bts.rds`
- Corrected input artifact: `analysis/outputs/family_sporc_corrected_full_age_bts_input_list.rds`
- Legacy adjusted-ISS artifact: `analysis/outputs/family_sporc.rds`
- OSA diagnostic runner: `analysis/run_family_osa_afscosa.R`
- OSA diagnostic artifact: `analysis/outputs/family_diagnostics/osa/sporc_corrected_full_age_bts_afscosa.rds`
- Retrospective runner: `analysis/run_family_retro_9.R`
- Retrospective data audit: `analysis/outputs/family_diagnostics/retro_9_data_availability.csv`
- Retrospective selectivity audit: `analysis/outputs/family_diagnostics/retro_9_selectivity_check.csv`
- afscOSA commit: 2dfe160f15b0d733d24305dd4f51e35a2fa60128
- SparseNUTS runner: `analysis/run_sporc_sparsenuts.R`
- SparseNUTS version: 1.0.2
- Fixed-steepness runner: `analysis/run_family_model_fixed_steepness.R`
- Fixed-steepness model artifact: `analysis/outputs/family_sporc_fixed_steepness.rds`
- Fixed-steepness tuned MCMC artifact: [fixed-steepness SparseNUTS result](#)
- Fixed-steepness MCMC source-fit MD5: a4a98d2ee88614aa9b8dc6489be5056e
- BTS parameter-RW runner: `analysis/run_family_model_bts_parameter_rw.R`
- Correlated BTS parameter-RW artifact: `analysis/outputs/family_sporc_fixed_steepness_bts_mvrw.rds`
- Correlated BTS enhancement base commit: 63af5a30dedeb38fbc2dea347748441c2f5f3af1
- Correlated BTS local-source changes present: TRUE
- Correlated BTS local-source fingerprint: 8aac0c131d6182122aa5d290b4e62934
- Correlated BTS parameter-RW MCMC artifact: [sparse-metric correlated-RW SparseNUTS result](#)
- Correlated BTS parameter-RW MCMC source-fit MD5: b7d4f30d323167a3f73cc2b6b1ce833e
- Projection input adapter: `R/write_spmr_projection_inputs.R`

- Projection runner: `analysis/run_sporc_spmr_projection.R`
- spmR version: 0.3.0
- spmR source commit: `e86fc6aa6f48a4de29d1deba0e8bd1df93b8717f`
- Projection source-fit MD5: `17288f8dd88bf766cdb3f5b7bce82068`
- ASAR version: 2.5.0.9000
- ASAR source commit: `e61c0f7d15569163fd5be8a7bfc582d5ca96c9af`
- ASAR-format project glossary: `reporting/report_glossary.tex`
- Primary model artifact: `analysis/outputs/family_sporc_corrected_full_age_bts.rds`
- Primary input and mapping artifact: `analysis/outputs/family_sporc_corrected_full_age_bts_input_list.rds`
- Model-family registry: `config/model_family.yml`

```
Rscript analysis/build_corrected_full_age_bts_data.R
SPORC_CASE_ID=sporc_corrected_full_age_bts SPORC_OUTPUT_STEM=family_sporc_corrected_full_age_bts SPORC_
Rscript analysis/run_family_osa_afscosa.R
Rscript analysis/run_family_retro_9.R
Rscript analysis/run_family_self_test_30.R
Rscript analysis/run_family_model_fixed_steepness.R
Rscript analysis/run_sporc_sparsenuts.R --fit-stem=family_sporc_fixed_steepness --model-name=sporc_corr
Rscript analysis/run_family_model_bts_parameter_rw.R
Rscript analysis/run_sporc_sparsenuts.R --fit-stem=family_sporc_fixed_steepness_bts_mvrv --model-name=s
Rscript analysis/run_sporc_spmr_projection.R
quarto render reporting/EBS_wp_sporc.qmd --to html
cp reporting/EBS_wp_sporc.html docs/index.html
```

Appendix

Acronyms

This appendix follows the [Alaska Stock Assessment Report \(ASAR\) package](#) glossary convention. The project glossary uses ASAR’s `\newacronym` format, retains shared package entries where available, and adds terms specific to the SPoRC bridge and its diagnostics.

Session information

R version 4.6.1 (2026-06-24)

Platform: aarch64-apple-darwin23

Running under: macOS Tahoe 26.6.1

Matrix products: default

BLAS: /Library/Frameworks/R.framework/Versions/4.6/Resources/lib/libRblas.0.dylib

LAPACK: /Library/Frameworks/R.framework/Versions/4.6/Resources/lib/libRlapack.dylib; LAPACK version 3.

locale:

[1] C.UTF-8/C.UTF-8/C.UTF-8/C/C.UTF-8/C.UTF-8

time zone: America/Los_Angeles

tzcode source: internal

attached base packages:

[1] stats graphics grDevices utils datasets methods base

other attached packages:

[1] yaml_2.3.12 gt_1.3.0 ggthemes_5.2.0 ggplot2_4.0.3

Table 23: Acronyms used in this working paper.

Acronym	Definition
AD	automatic differentiation
ADMB	AD Model Builder
ASAR	Alaska Stock Assessment Report
ATS	acoustic-trawl survey
AVO	acoustic vessels of opportunity
BTS	bottom trawl survey
CPUE	catch per unit of effort
EBS	eastern Bering Sea
ESS	effective sample size
HMC	Hamiltonian Monte Carlo
ISS	input sample size
MCMC	Markov Chain Monte Carlo analysis
MLE	maximum-likelihood estimate
NUTS	No-U-Turn sampler
OSA	one-step-ahead
R-hat	rank-normalized split-chain convergence statistic
RMS	root mean square
SDNR	standard deviation of normalized residuals
SPoRC	Stock Population and Recruitment Computational framework
SSB	spawning stock biomass

[5] asar_2.5.0.9000 spmR_0.3.0 SPoRC_1.2.0.9000 RTMB_1.9

loaded via a namespace (and not attached):

[1] tidyselect_1.2.1	dplyr_1.2.1	farver_2.1.2
[4] loo_2.10.0	S7_0.2.2	fastmap_1.2.0
[7] PBSmodelling_2.70.2	XML_3.99-0.23	digest_0.6.39
[10] timechange_0.4.0	lifecycle_1.0.5	StanHeaders_2.32.10
[13] magrittr_2.0.5	compiler_4.6.1	rlang_1.3.0
[16] tools_4.6.1	knitr_1.51	labeling_0.4.3
[19] pkgbuild_1.4.8	curl_7.1.0	plyr_1.8.9
[22] xml2_1.6.0	RColorBrewer_1.1-3	abind_1.4-8
[25] afscOSA_0.0.1	withr_3.0.3	purrr_1.2.2
[28] grid_4.6.1	stats4_4.6.1	diffobj_0.3.6
[31] future_1.70.0	inline_0.3.21	progressr_1.0.0
[34] globals_0.19.1	scales_1.4.0	MASS_7.3-65
[37] tinytex_0.60	cli_3.6.6	rmarkdown_2.31
[40] crayon_1.5.3	generics_0.1.4	otel_0.2.0
[43] RcppParallel_5.1.11-2	future.apply_1.20.2	reshape2_1.4.5
[46] tzdb_0.5.0	commonmark_2.0.0	rstan_2.32.7
[49] stringr_1.6.0	tmbstan_1.1.0	parallel_4.6.1
[52] matrixStats_1.5.0	vctrs_0.7.3	V8_8.2.0
[55] Matrix_1.7-5	jsonlite_2.0.0	litedown_0.9
[58] hms_1.1.4	patchwork_1.3.2	listenv_1.0.0
[61] tidyr_1.3.2	glue_1.8.1	parallelly_1.48.0
[64] randtests_1.0.2	codetools_0.2-20	cowplot_1.2.0
[67] lubridate_1.9.5	stringi_1.8.7	gtable_0.3.6
[70] QuickJSR_1.10.0	tibble_3.3.1	pillar_1.11.1
[73] htmltools_0.5.9	R6_2.6.1	TMB_1.9.21
[76] tcltk_4.6.1	evaluate_1.0.5	lattice_0.22-9
[79] markdown_2.0	readr_2.2.0	png_0.1-9
[82] Rcpp_1.1.2	gridExtra_2.3.1	xfun_0.60
[85] fs_2.1.0	pkgconfig_2.0.3	