NC STATE UNIVERSITY

Modeling spatial distributions and habitat for species off the southeastern US coast

Reef fishes: distribution models, indicators and habitat from the SERFS trap–video survey, 1990–2024

Jie Cao, Kevin Craig, Janet Nye, Jade Hong

North Carolina State University · NOAA Southeast Fisheries Science Center · UNC Institute of Marine Sciences

SAFMC contract SA-25-111

SAFMC Scientific and Statistical Committee · 21 October 2026

The project: SA-25-111

Modeling spatial distributions and habitat for species off the southeastern US coast

SAFMC contract under NOAA's Inflation Reduction Act climate award

February 2025 – June 2027

Team

- Jie Cao, Kevin Craig — reef fishes
- Janet Nye, Jade Hong — pelagics
- Project oversight team
- Reviews: Habitat AP, SSC, Council

Four objectives

1. Spatiotemporal models integrating survey data
2. Habitat and environmental drivers of distribution
3. Predicted habitat areas vs current EFH
4. Abundance indices and range-shift metrics

Two species groups, two modeling tracks

This talk: the reef-fish track. Jade Hong presents the pelagic track.

Reef fishes — NCSU

- Snapper-grouper complex, 55 managed species
- SERFS trap–video survey, 1990–2024
- Single-species and joint-species models

Pelagic species — UNC IMS

- Coastal migratory pelagics, dolphin–wahoo
- Fishery-independent and -dependent data
- Single-species models



NC STATE UNIVERSITY

This talk

Preliminary results — the aim is SSC feedback on the modeling approach

1. **Challenges**
   Uneven data across species; two gears with different histories and catchability

2. **Modeling framework**
   Two-gear integration; single-species and multispecies structures

3. **Simulation testing**
   Can the estimator recover a known truth at survey sample sizes?

4. **Diagnostics**
   AIC in sample; three cross-validation designs out of sample

5. **Sensitivity runs**
   Era (2010 onward), gear (trap only), covariate choice

6. **Preliminary products**
   Indices, distribution indicators, density maps, EFH tiers

One survey, two gears, two histories

SERFS chevron trap since 1990; underwater video mounted on the trap since 2010



~43,000
paired samples
(one species' data)

~25,000
locations, Cape Hatteras
to Florida

34
survey years
(none in 2020)

2 gears
same drop,
different detection

The challenge. The two gears detect species differently, and only the trap covers 1990–2009. A model that combines them must separate the population signal from the gear signal.

NC STATE UNIVERSITY

Why combine the two gears?

Two observation processes, one population

Coverage and complementarity

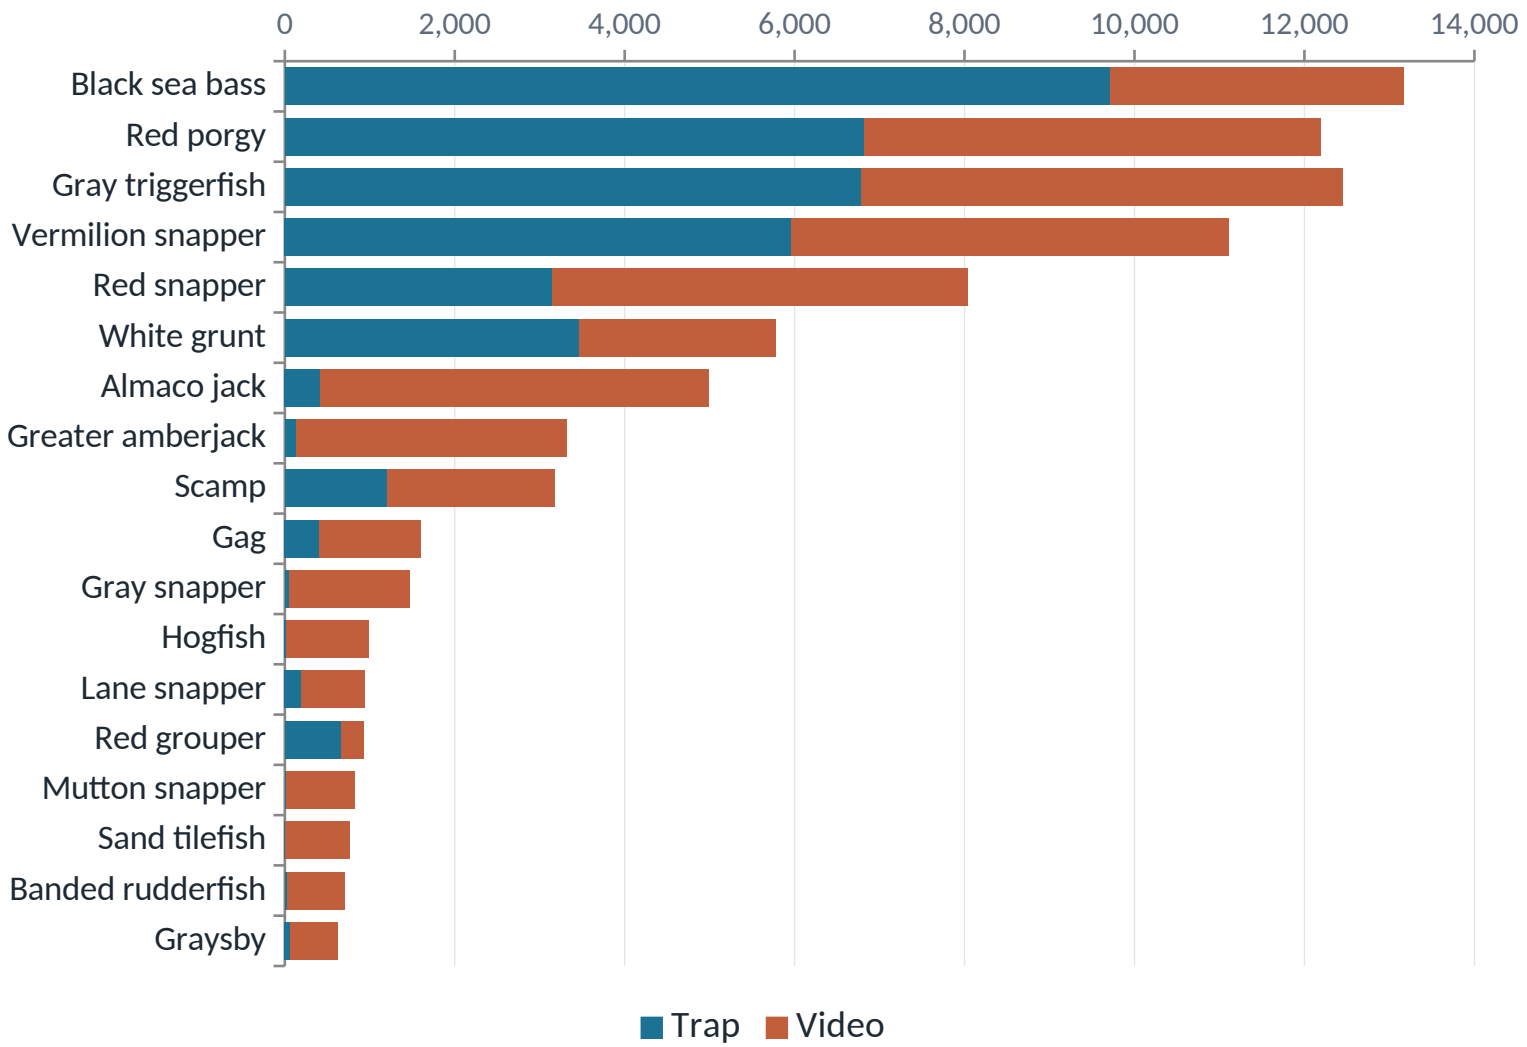
- Expand spatial and temporal coverage beyond either survey
- Reduce dependence on the limitations of any single gear
- Use complementary observation processes to characterize distributions across the shelf
- Separate the shared population pattern from trap- and video-specific variability

One index, 1990–2024

- Trap series extends the population signal to the years before video (pre-2010)
- Video strengthens estimation where the surveys overlap (2010–2024)
- A single annual index from the shared spatiotemporal component
- Better continuity, and potentially precision, than either survey alone

The data are uneven across species

Positive samples, trap and video combined, 1990–2024: the 18 species in the multispecies model



Six species

≥ 2,300 positives in each gear, all 34 years: their own models

Why 18

≥ 600 positives combined; four video-only

Sharing information

Species share habitat. A few community patterns are estimated jointly; a rare species' map is a weighted blend of them — its data only set the weights.

Two frameworks, one estimation engine

Both fitted in tinyVAST (Thorson et al. 2024): SPDE spatial fields, Laplace approximation, TMB

Part I — single-species (6 species)

- Each species on its own data
- Fields shared by both gears, with a trap loading
- Gear-specific residual fields
- Covariates chosen by AIC, checked out of sample

Part II — multispecies (18 species)

- Species share a few community factors
- Own loadings and own residual field per species
- Data-limited species borrow spatial structure
- Same two-gear treatment

How they relate. The single-species models are the standard. The multispecies model must reproduce their maps before its maps for poorly sampled species are trusted.

Part I

Single-species models

Six species · structure, checks, sensitivity, products

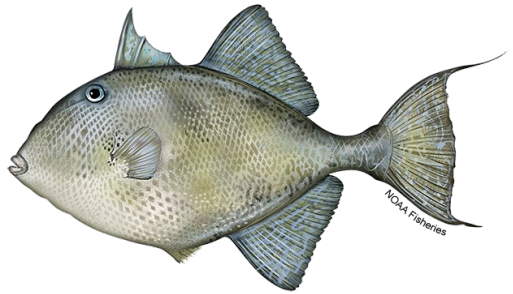
NC STATE UNIVERSITY

The six species with their own models

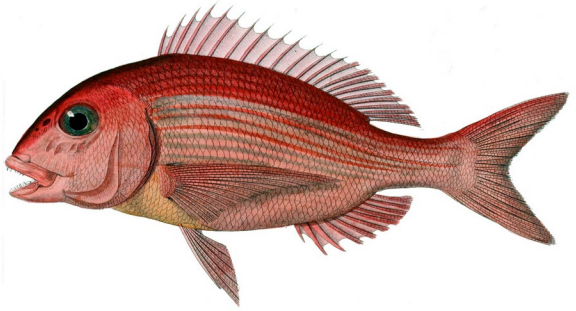
Thousands of positive samples in each gear, detected in all 34 survey years



Black sea bass
Centropristis striata



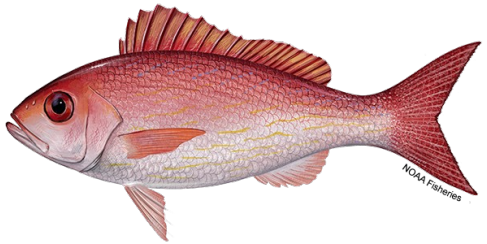
Gray triggerfish
Balistes capriscus



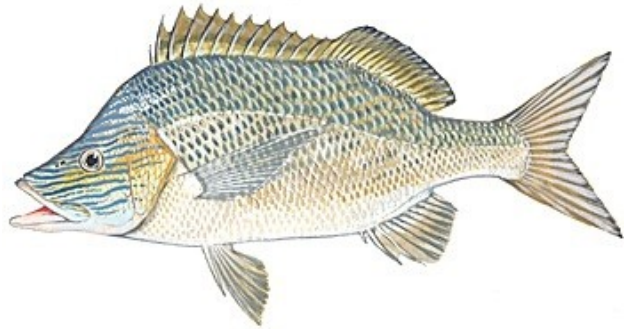
Red porgy
Pagrus pagrus



Red snapper
Lutjanus campechanus



Vermilion snapper
Rhomboplites aurorubens



White grunt
Haemulon plumierii

What a spatiotemporal model does

The expected catch at a site in a year is built from a few interpretable pieces

Expected catch = **Persistent map** where the species lives on average × **Annual map** how that pattern shifts this year × **Habitat** depth and bottom temperature × **Year level** overall abundance this year × **Gear and effort** what this gear catches per drop

Persistent spatial field

- Constant geographic structure: enduring differences in habitat, bathymetry, substrate or survey availability
- Does not change from year to year; one surface for 1990–2024

Spatiotemporal field

- Annual departures from the persistent pattern: redistribution, recruitment pulses, episodic conditions, interannual movement
- Estimated separately for each year; 34 surfaces

Common approaches to integrating surveys

How the literature combines multi-gear surveys

1. **Survey-specific catchability intercepts**
   A constant gear effect as a scalar offset — the simplest and most common

2. **Catchability covariates**
   Gear efficiency varies with environmental or habitat covariates

3. **Spatially varying catchability**
   Gear efficiency varies continuously across space

Observation-scale prediction

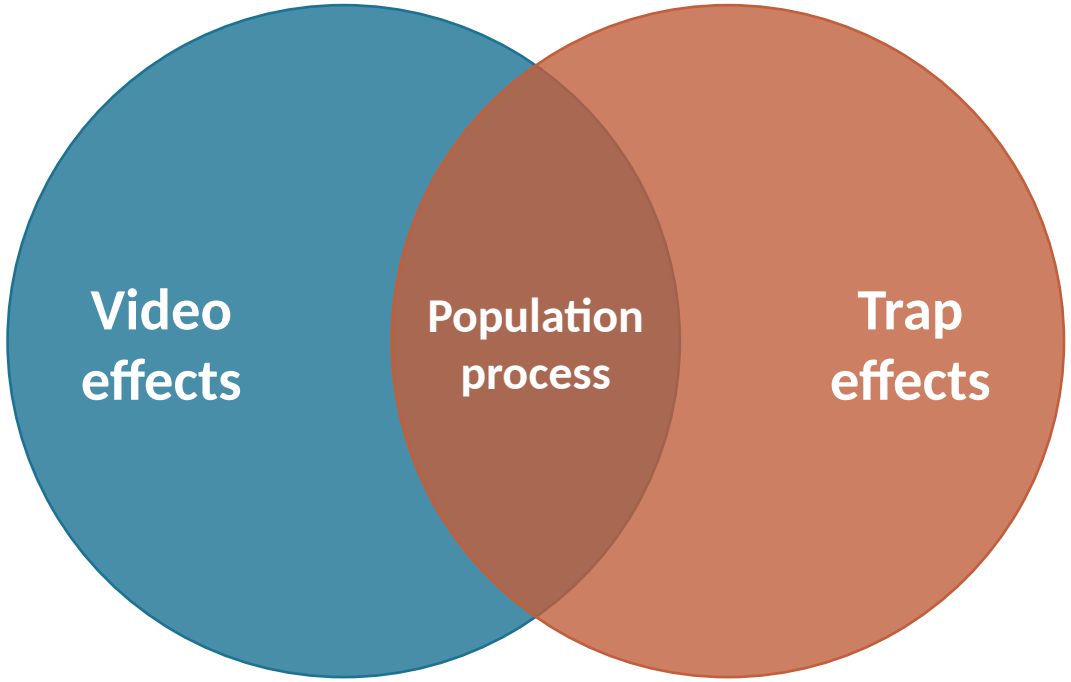
The population process combined with one survey's gear effects — what these approaches estimate

Population-scale prediction

The underlying density with gear effects removed — the quantity this study sets out to estimate

Separating survey effects from the population signal

Existing approaches are survey-conditional; paired trap and video samples let us do better



Video observations

shared process + video catchability, selectivity and residual spatial variation

Shared population process

the spatial and temporal patterns both surveys support at once

Trap observations

shared process + trap catchability, saturation and residual spatial variation

What each survey actually sees

Each gear has its own persistent field and its own annual field — how much of them is the population?

Trap lens

- Persistent field ω_trap(s): where the trap catches this species, year after year
- Annual field ε_trap(s,t): how the trap's pattern shifts each year

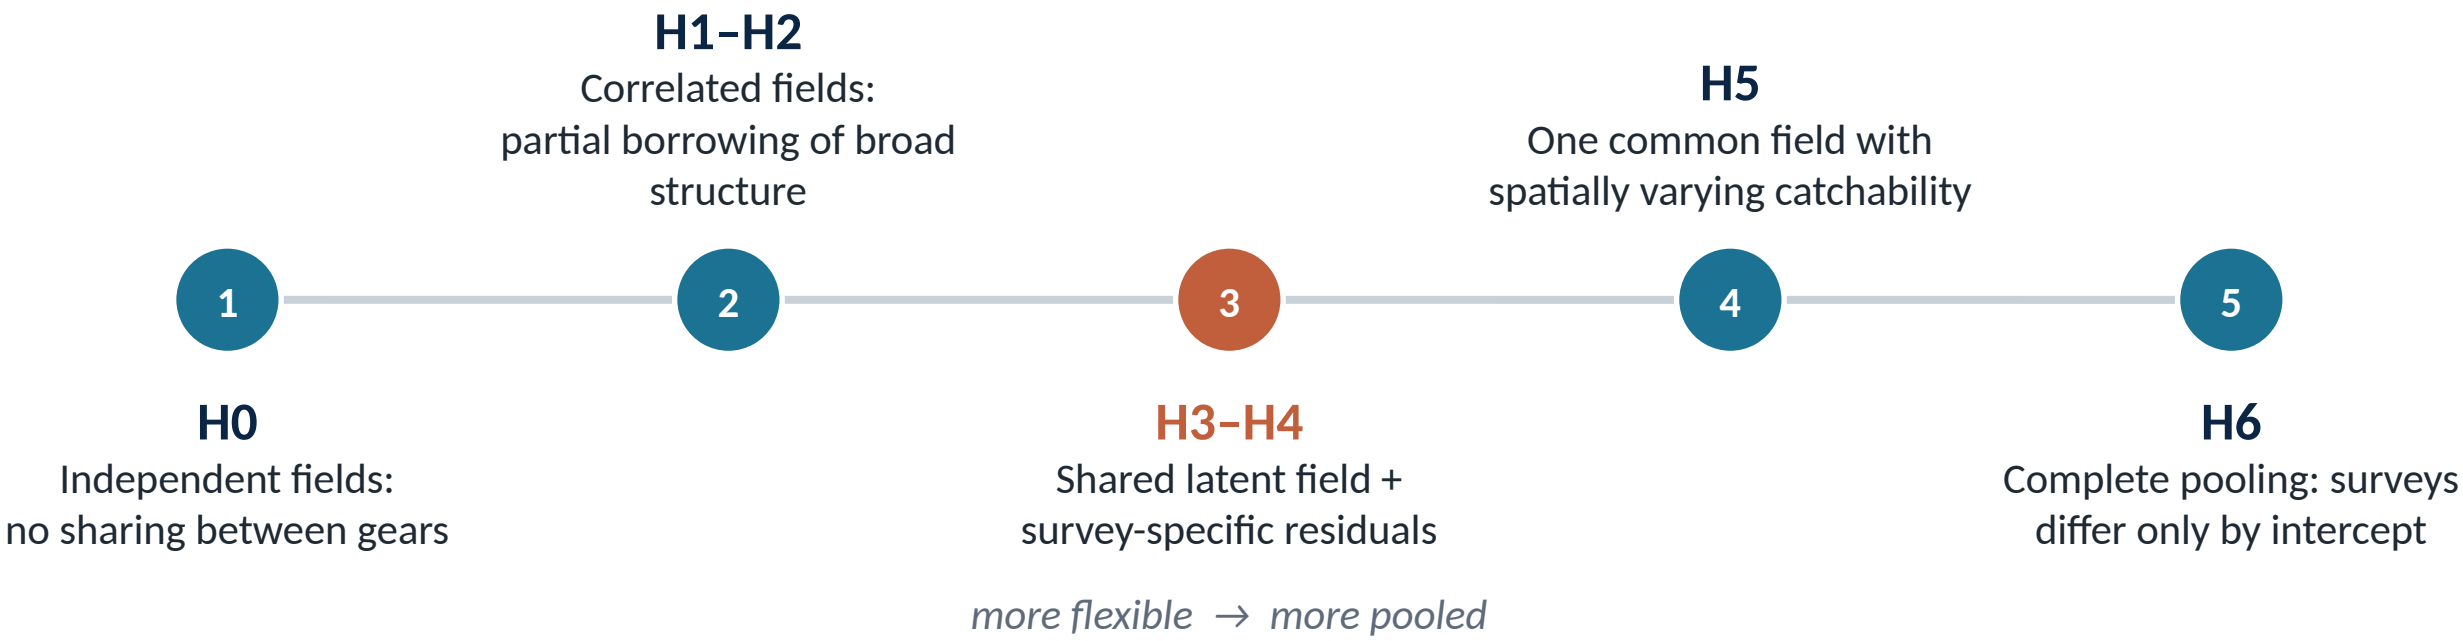
Video lens

- Persistent field ω_video(s): where the video sees this species, year after year
- Annual field ε_video(s,t): how the video's pattern shifts each year

The question. How much of each field is population signal shared by both gears, and how much is specific to one gear? Every answer is a different model — nothing shared, everything shared, or something in between.

Seven hypotheses

H0–H6 span a continuum from independent survey fields to complete pooling



Focus of this work. H4 decomposes each survey into a shared population process plus a survey-specific residual field — the only model that yields a population-scale estimate distinct from either gear alone.

The seven models, side by side

Spatial structure, spatiotemporal structure, survey effects and the resulting estimator

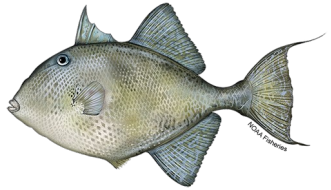
Model	Spatial	Spatiotemporal	Survey effects	Estimator
H0	separate fields	separate fields	intercepts; fully separate fields	separate trap and video indices
H1	separate, estimated correlation	separate fields	broad structure may be shared; annual independent	separate indices
H2	as H1	common signal, fixed loading	temporal change shared more strongly	separate indices
H3	as H1	separate, estimated correlation	sharing estimated for both components	separate indices
H4	**shared field + loadings + residuals**	**shared field + loadings + residuals**	**common signal separated from survey pattern**	**index from the shared population component**
H5	one common field	one common field	catchability varies spatially (SVC)	common field net of SVC
H6	one common field	one common field	surveys differ only by intercept	fully pooled field

Fitting the seven models to real data

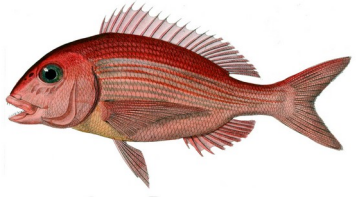
All seven configurations fitted to each of the six species on identical data, mesh and covariates



Black sea bass



Gray triggerfish



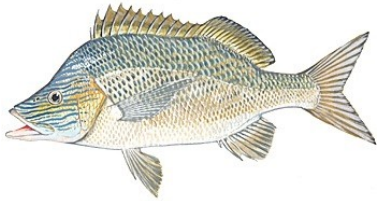
Red porgy



Red snapper



Vermilion snapper



White grunt

7 models × 6 species = 42 fits compared by AIC within species

- **Which gear-sharing assumptions does the survey itself support?**
- **Does the shared-field model (H4) outperform the alternatives on real data?**

Relative support for the seven models

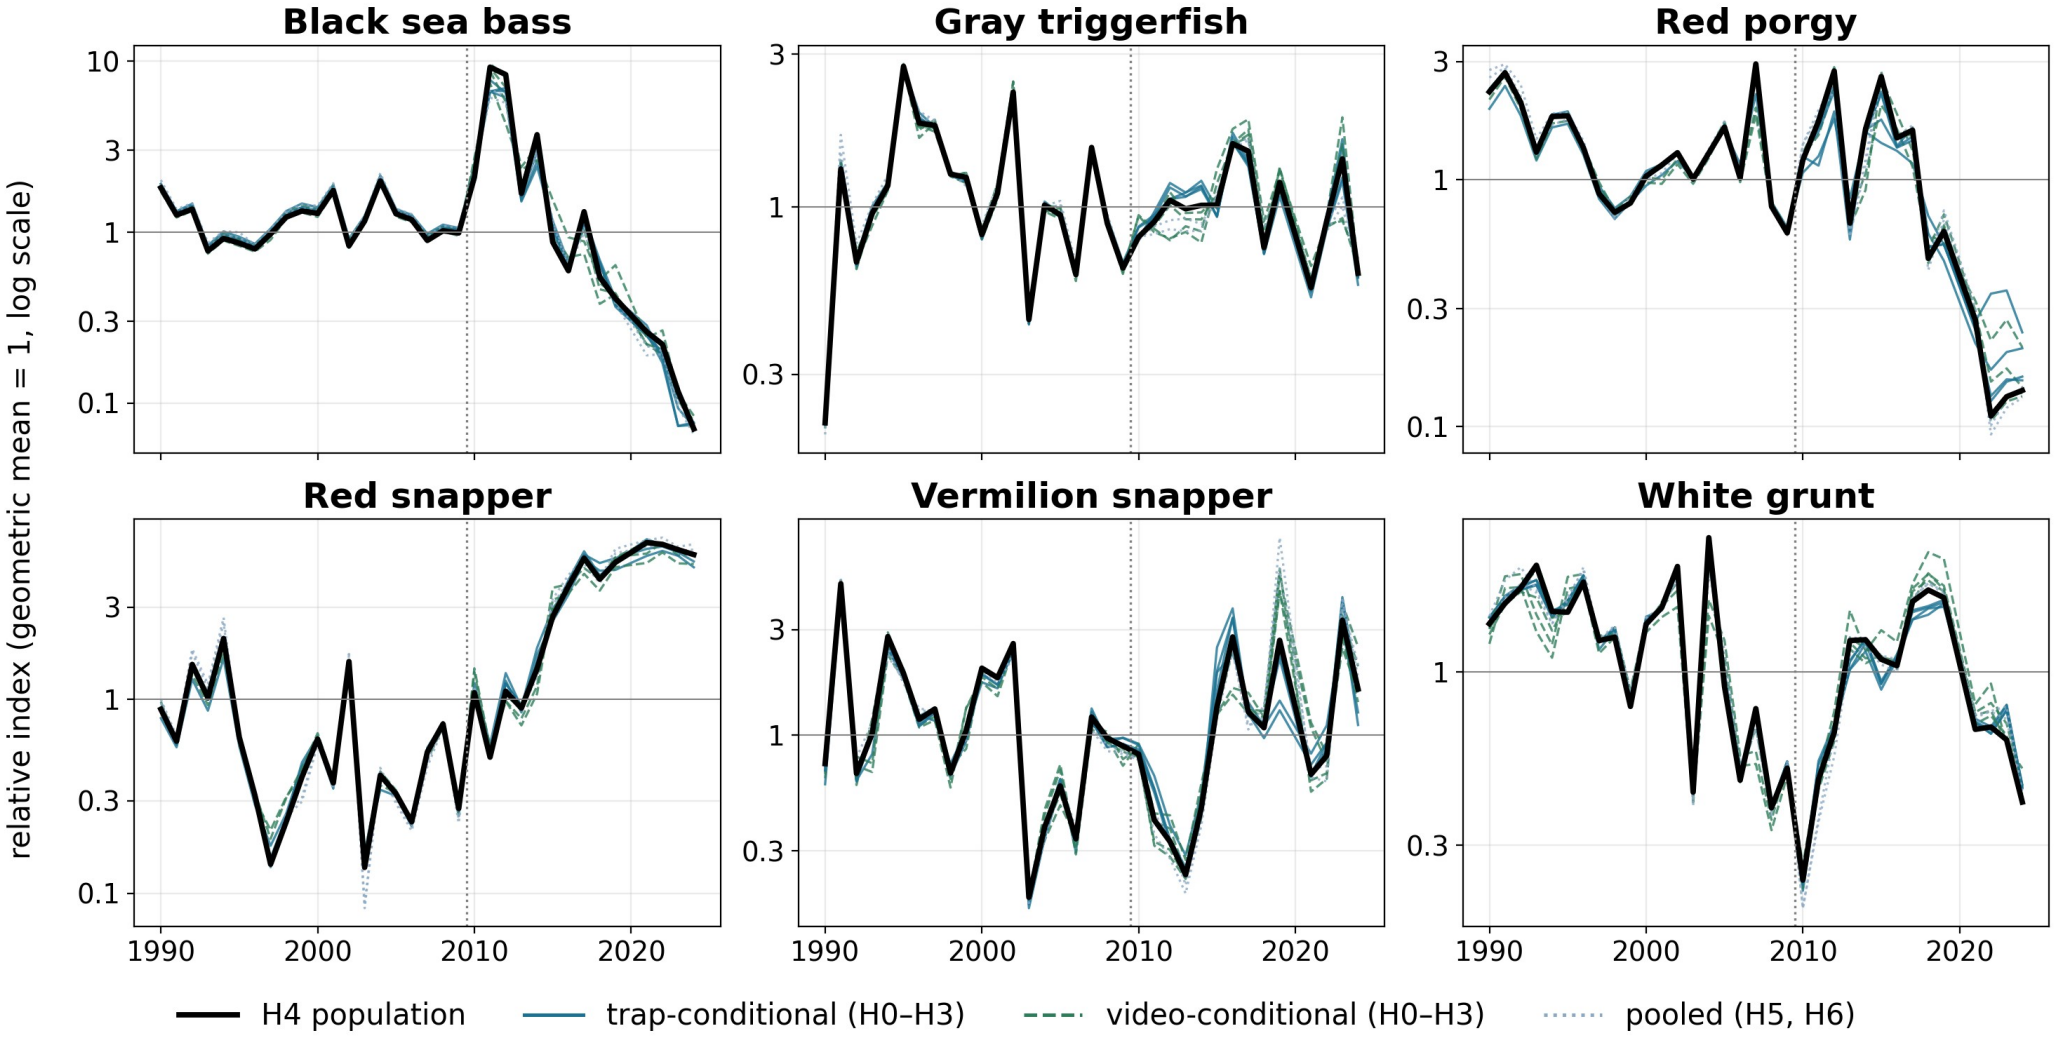
ΔAIC from the lowest AIC for each species; bold marks the best-supported models

Species	H0	H1	H2	H3	H4	H5	H6
Black sea bass	737	390	48	**0**	**0**	1,908	2,481
Gray triggerfish	678	387	13	**0**	**0**	4,278	4,636
Red porgy	702	314	6	**0**	**0**	4,480	4,722
Red snapper	722	396	6	**0**	**0**	1,398	1,445
Vermilion snapper	694	482	31	**0**	**0**	4,195	4,652
White grunt	327	170	16	**0**	**0**	2,536	2,795

Reading. Pooling (H5, H6) is rejected by thousands of AIC units: the gears do not see the same map. Full independence (H0) is rejected too. H3 and H4 tie because H4 is H3 re-expressed with an explicit shared component; H4 is preferred because it defines the population-scale estimand.

Annual trends agree across estimators

Eleven estimators from the seven models, each standardised to a mean of 1; dashed line = 2010



NC STATE UNIVERSITY

Stress-testing the models against a known truth

Simulation: the fitted H4 shared field is the truth; new trap and video data are generated from it

Fit H4 to each species → **Take its shared field** as the true population → **Simulate** trap and video data → **Refit all 7 models** to the simulated data → **Score** each estimator against truth

Residual variation

how much the gears disagree beyond the shared signal

0.25×	1×	2×
Low	Moderate	High

Trap compression

a weaker trap echo of the shared signal

1.00	0.75	0.50
None	Mild	Strong

Video duration

longer or shorter overlap between surveys

2000	2010	2015
Long	Current	Short

9 scenarios

6 species

7 models per fit

100 replicates

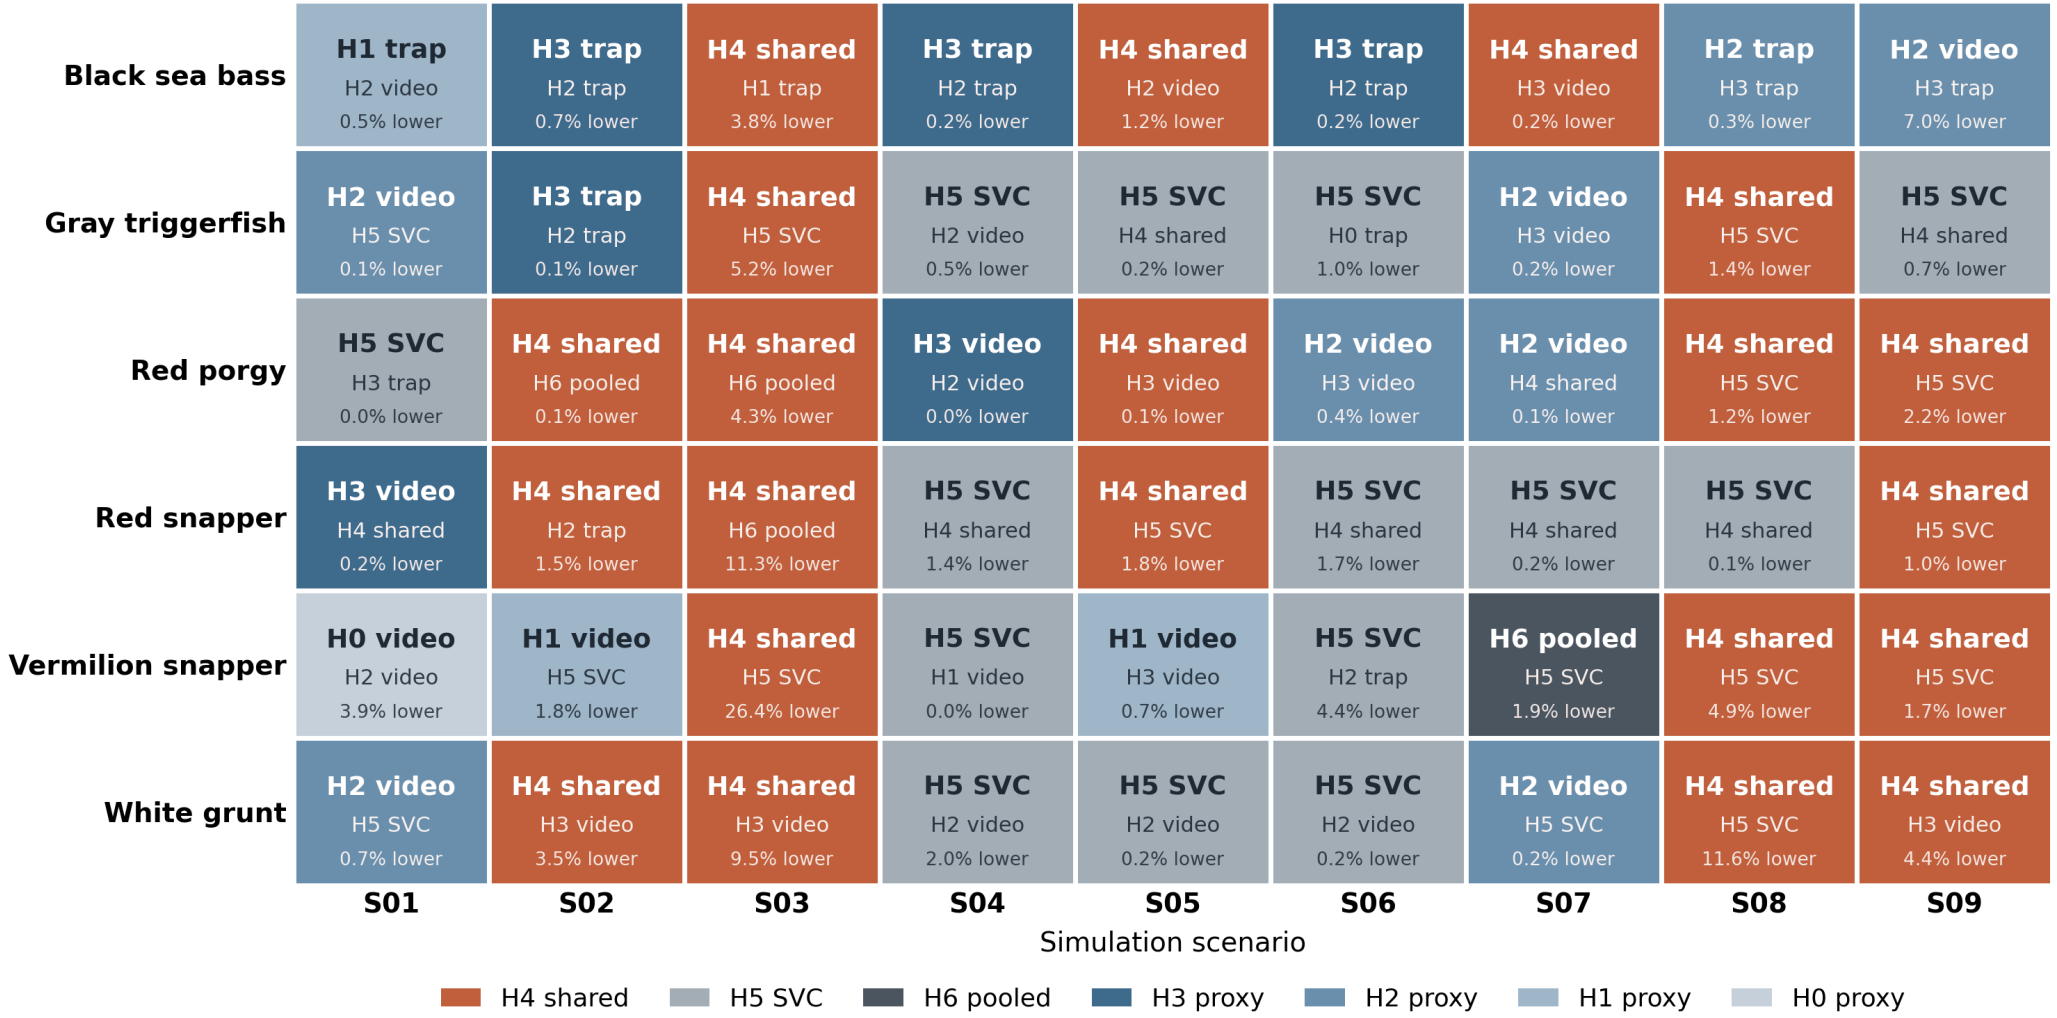
37,800 model fits

The nine scenarios

Scenario	Residual level	Multiplier	Trap loading	Video series	Start year
S01	Low	0.25	Equal loading	Current	2010
S02	Moderate	1.00	Equal loading	Current	2010
S03	High	2.00	Equal loading	Current	2010
S04	Moderate	1.00	Mild compression	Current	2010
S05	Moderate	1.00	Strong compression	Current	2010
S06	Moderate	1.00	Mild compression	Long	2000
S07	Moderate	1.00	Mild compression	Short	2015
S08	High	2.00	Mild compression	Long	2000
S09	High	2.00	Mild compression	Short	2015

Simulation result 1 — annual index

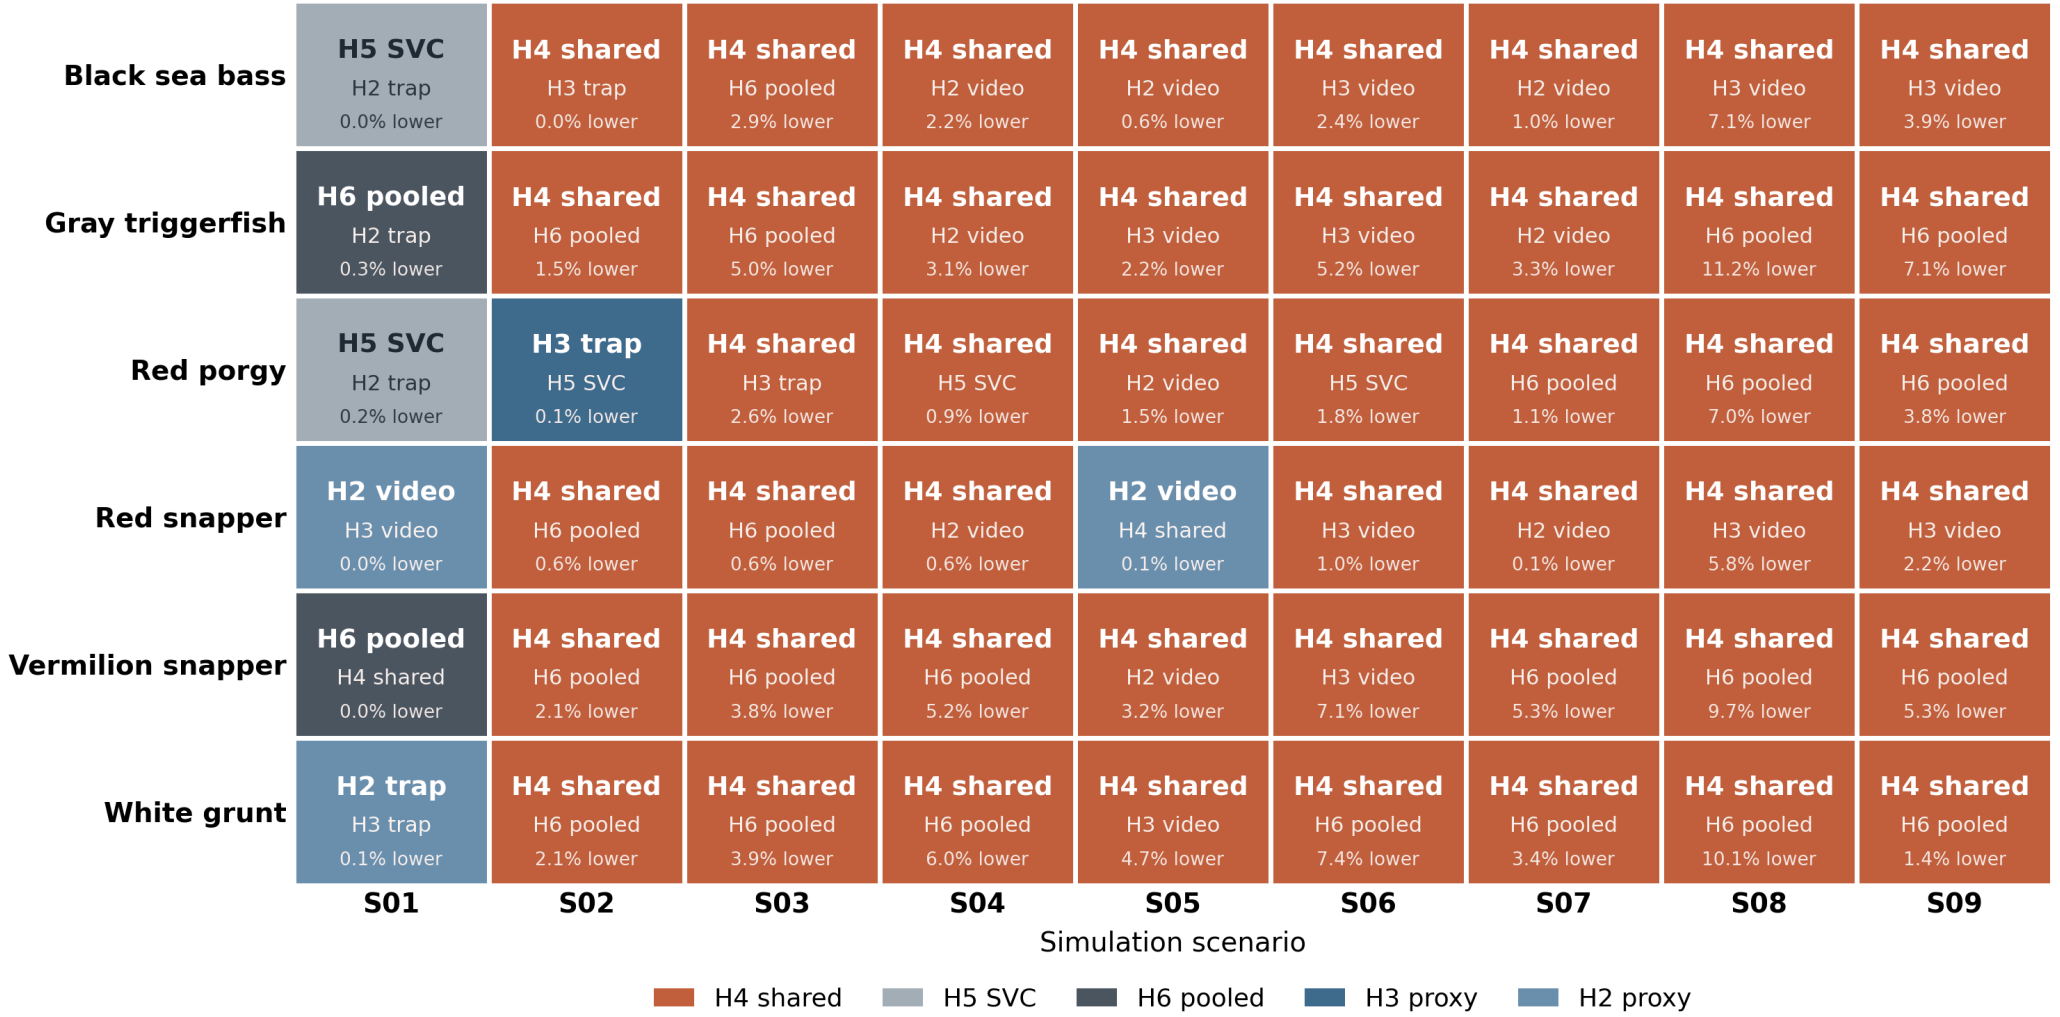
Best estimator by median RMSE for each species and scenario; smaller text gives the runner-up and its RMSE gap



NC STATE UNIVERSITY

Simulation result 2 — population field

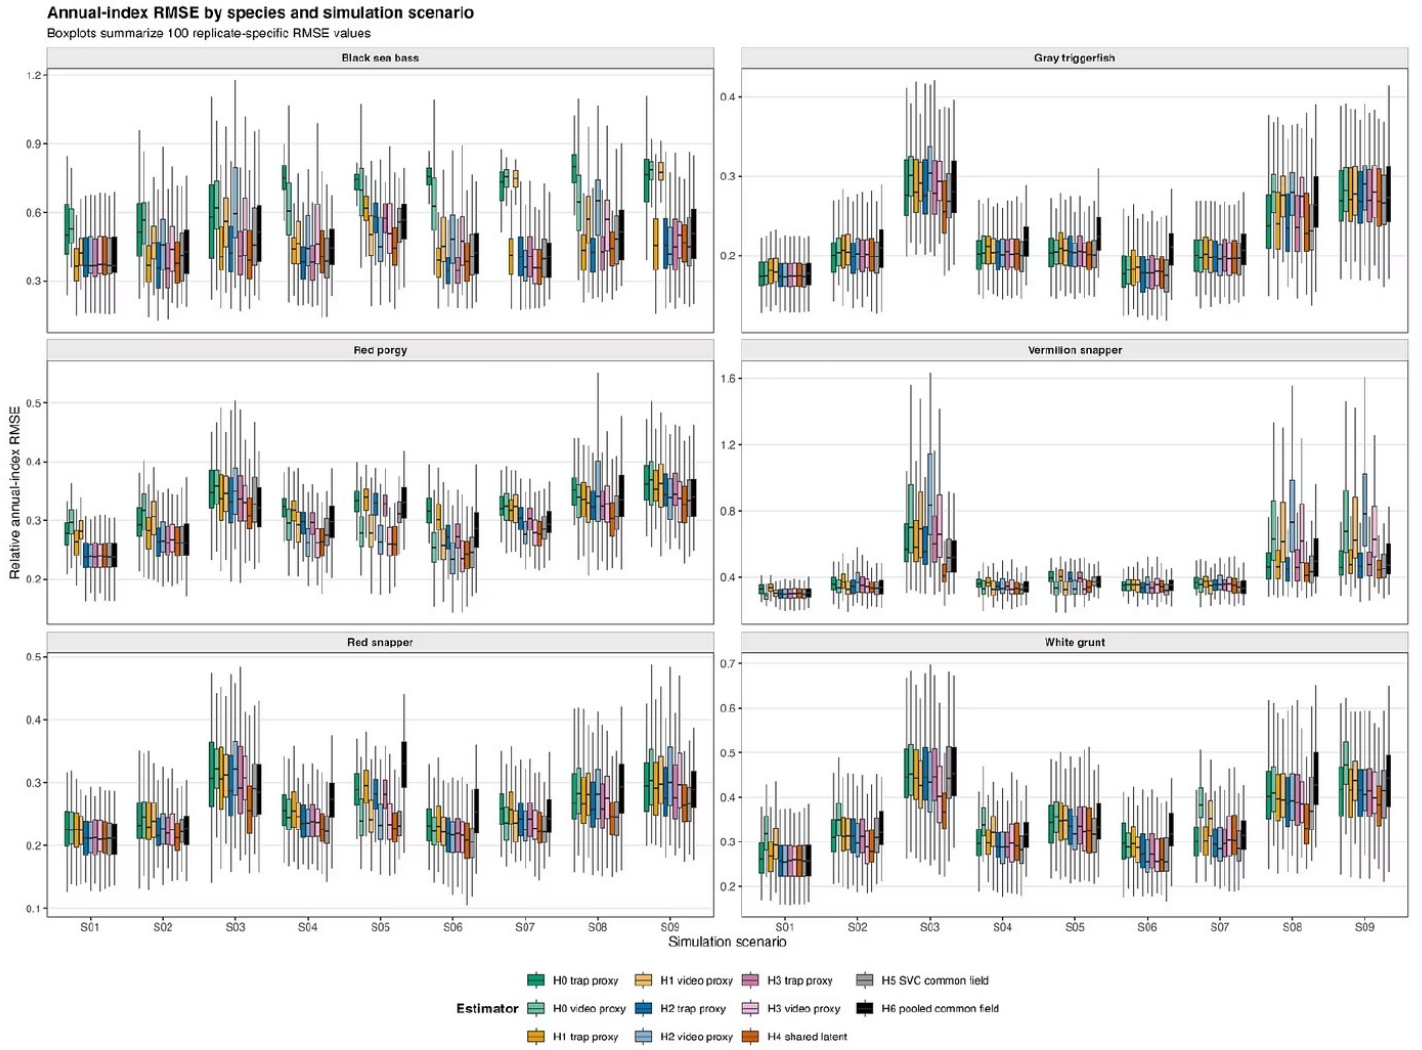
Best estimator of the latent population field by median RMSE; smaller text gives the runner-up and its RMSE gap



NC STATE UNIVERSITY

Simulation result 3 — robustness across replicates

Annual-index RMSE, 100 replicates per scenario, all eleven estimators



NC STATE UNIVERSITY

What the real data and the simulation tell us

On real data (AIC, six species)

- Complete pooling (H5, H6) rejected by thousands of AIC units: the gears do not see the same map
- Fully independent fields (H0) rejected by hundreds: the gears do share structure
- H3 and H4 tie — the data support a shared component plus gear-specific residuals
- All eleven estimators give the same annual trends; H4 runs through the middle

Against a known truth (37,800 fits)

- Annual index: H4 or H5 is the best estimator in most scenarios; H4 dominates once the trap response is compressed
- Population field: H4 is best in almost every species–scenario cell
- Robust across all 100 replicates, not only on average
- Single-gear proxies degrade under high residual variation and short overlap

Conclusion. H4 — a shared population field with gear loadings and gear-specific residuals — is the structure carried forward.

Comparing with a published non-spatial model

Conn (2010, Can. J. Fish. Aquat. Sci.) combines noisy CPUE indices with the same shared-signal + index-specific-noise logic

Conn (2010) — multi-index hierarchical model

- One shared log-abundance trend across all indices
- Index-specific scaling (catchability) and process error
- Known sampling error per index

H4 — the spatial extension

- Shared population process: the trend, but over space and time
- Survey loadings: the scaling, estimated per gear
- Survey residual fields: the process error, spatially structured

Same idea, same finding. Conn found the hierarchical approach helped most when indices disagreed strongly; our simulations show exactly the same pattern for H4.

Fitting choices, and how each was settled

Every choice was tested rather than assumed; the first row is what we have covered so far

Choice	Compared	Settled on	How
Gear structure	independent · one shared · shared + residual	**shared + residual (H4)**	AIC; simulation; cross-gear prediction
Mesh resolution	5, 10, 15 km	5 km (810 vertices)	AIC
Observation model	Poisson · NB · Tweedie	negative binomial	AIC; calibration
Year effects	annual field only · fixed effects	fixed effects + field	index identifiability
Covariates	depth, temperature: none / linear / smooth	species-specific	AIC, then cross-validation
Uncertainty	fixed + random · random only	random effects only	interval inflation

Objective 2

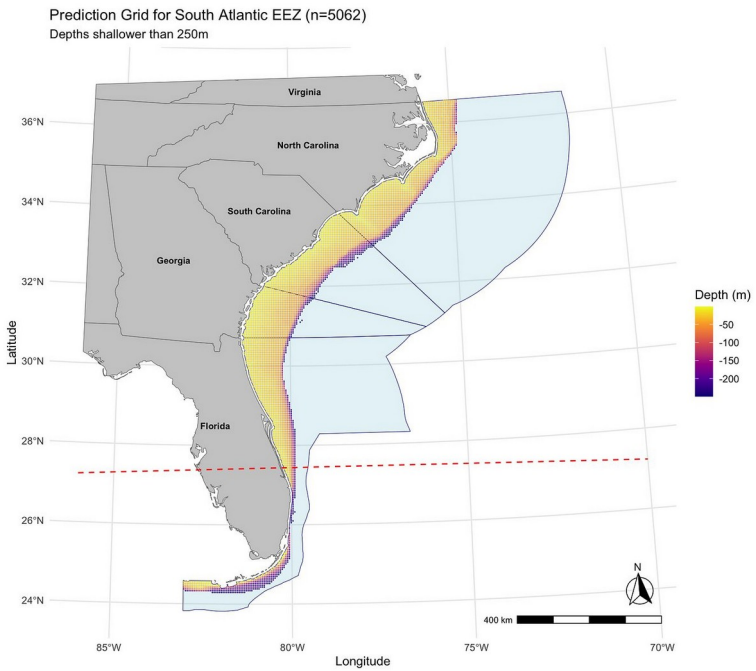
Evaluate habitat drivers

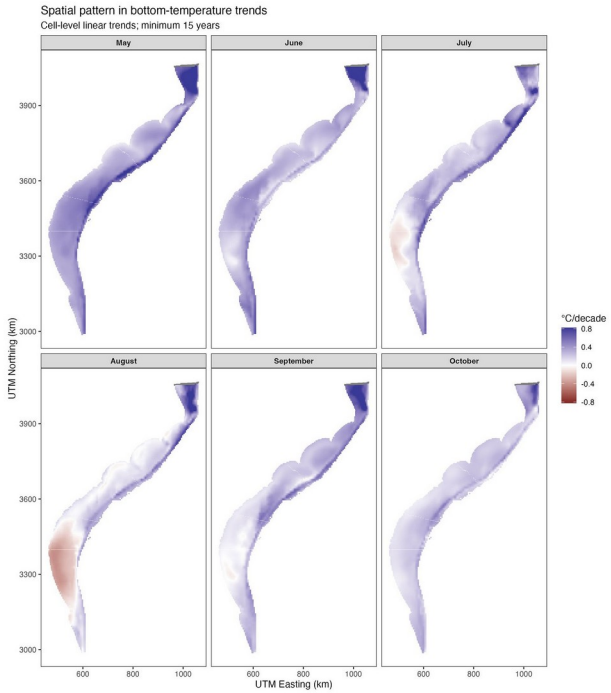
Quantify the relative importance of habitat features and environmental conditions in shaping species distributions

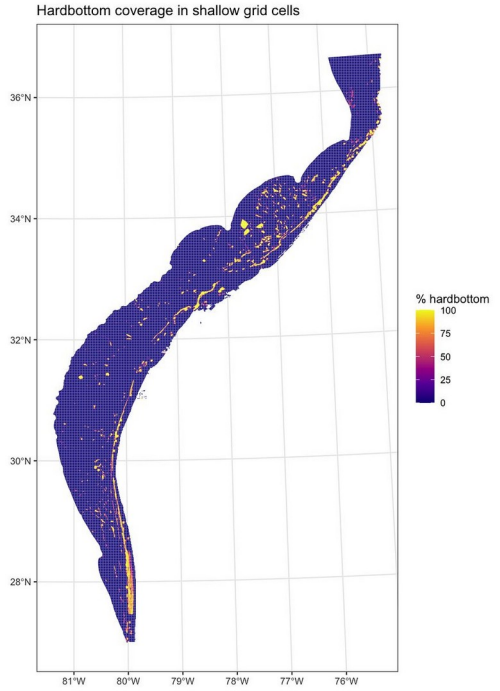
NC STATE UNIVERSITY

Covariates: depth, bottom temperature, hard bottom

One fixed habitat gradient, one climate variable, one substrate layer







Depth

Bathymetry on the 5-km prediction grid, shallower than 250 m

Bottom temperature

GLORYS12V1 reanalysis (CMEMS, 1/12°, 50 levels, 1993 onward), matched to each sample by month

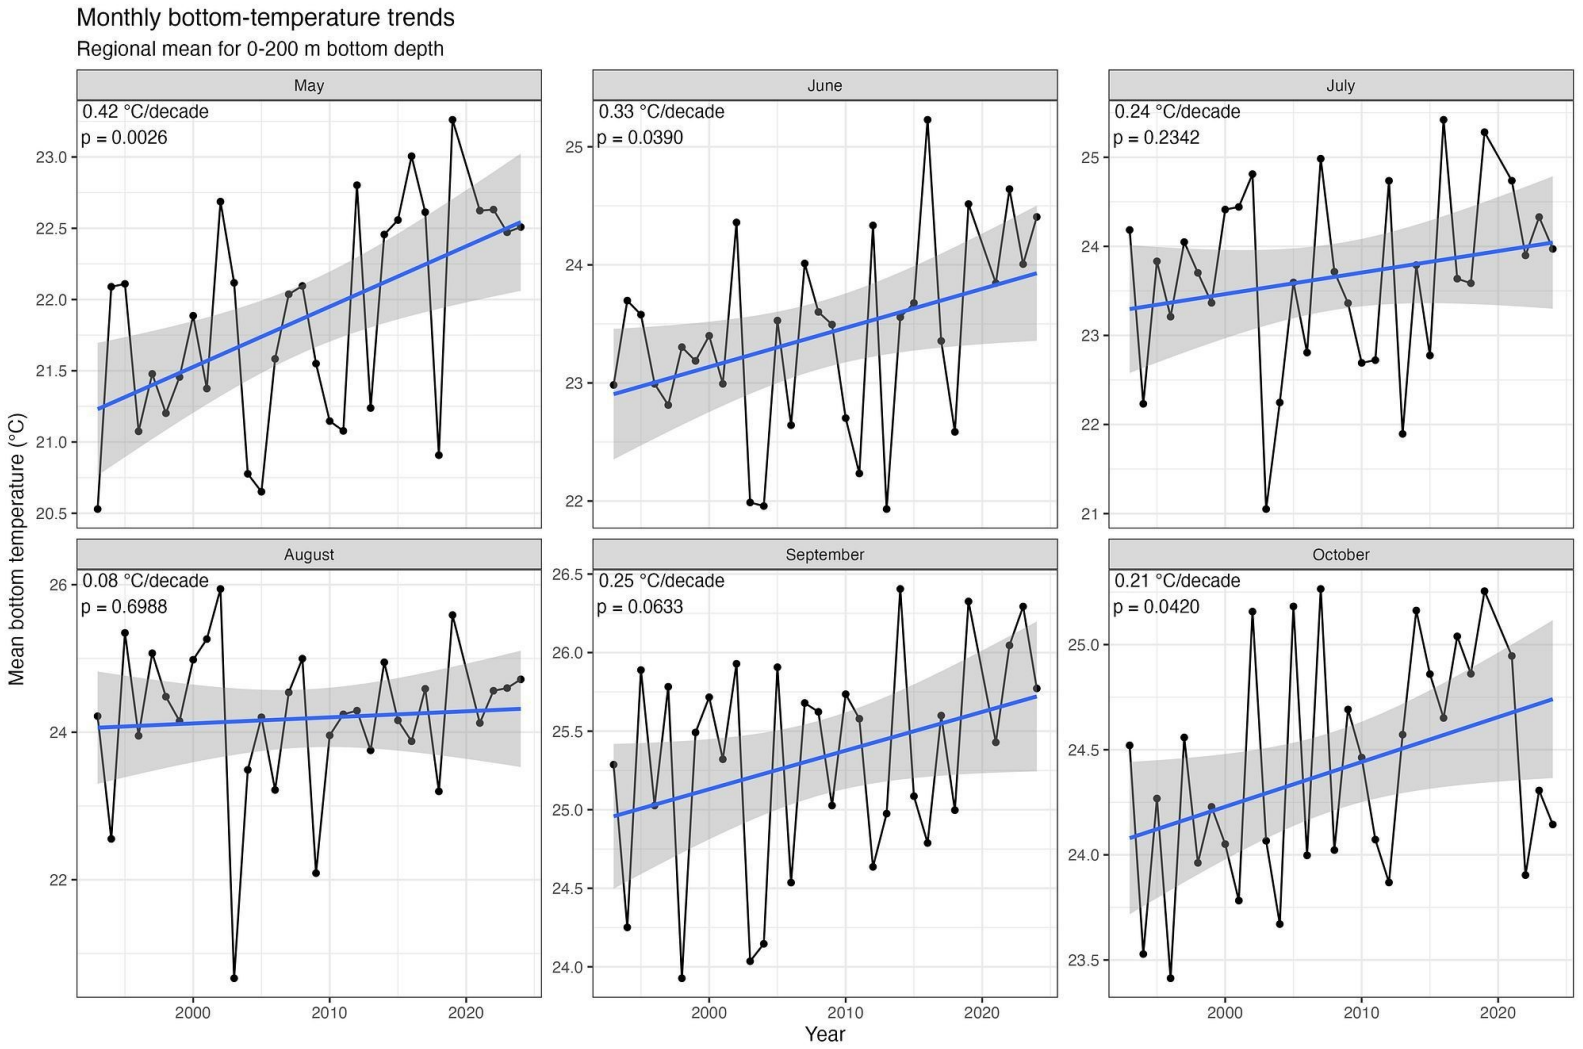
Hard bottom

Share of each cell covered by mapped hard-bottom polygons

NC STATE UNIVERSITY

Bottom temperature is rising across the shelf

GLORYS12 bottom temperature, regional mean for 0–200 m, by survey month



NC STATE UNIVERSITY

Covariate configurations: a full factorial

Every combination fitted — no stepwise search — so each covariate can be isolated later

Depth

absent · linear · spline

Temperature

absent · linear · spline

Hard bottom

absent · linear

18

covariate combinations
3 depth × 3 temperature × 2 hard bottom

6

species

108

model fits, identical mesh
and error distribution

Shared by all fits: H4 gear structure, survey-specific intercepts, year as a factor, log sampling-duration offset, thin-plate splines with k = 4.

Model selection: two criteria

AIC scores the fit; a second criterion checks whether the covariate took structure from the random fields

Criterion 1 — marginal AIC

- Fit to the counts, penalised for fixed-effect and variance parameters
- Random fields integrated out
- Cannot see whether a covariate took spatial structure away from the field or merely improved fit alongside it

Criterion 2 — process variance explained

- Variance of the predicted field (persistent and annual, separately) relative to the null model
- Evaluated at sampled locations within the species' observed depth range
- Positive: covariates absorbed structure the field would otherwise carry

The confound. The random fields and the covariates compete for the same spatial variance. AIC scores the outcome, not the division — hence the second criterion.

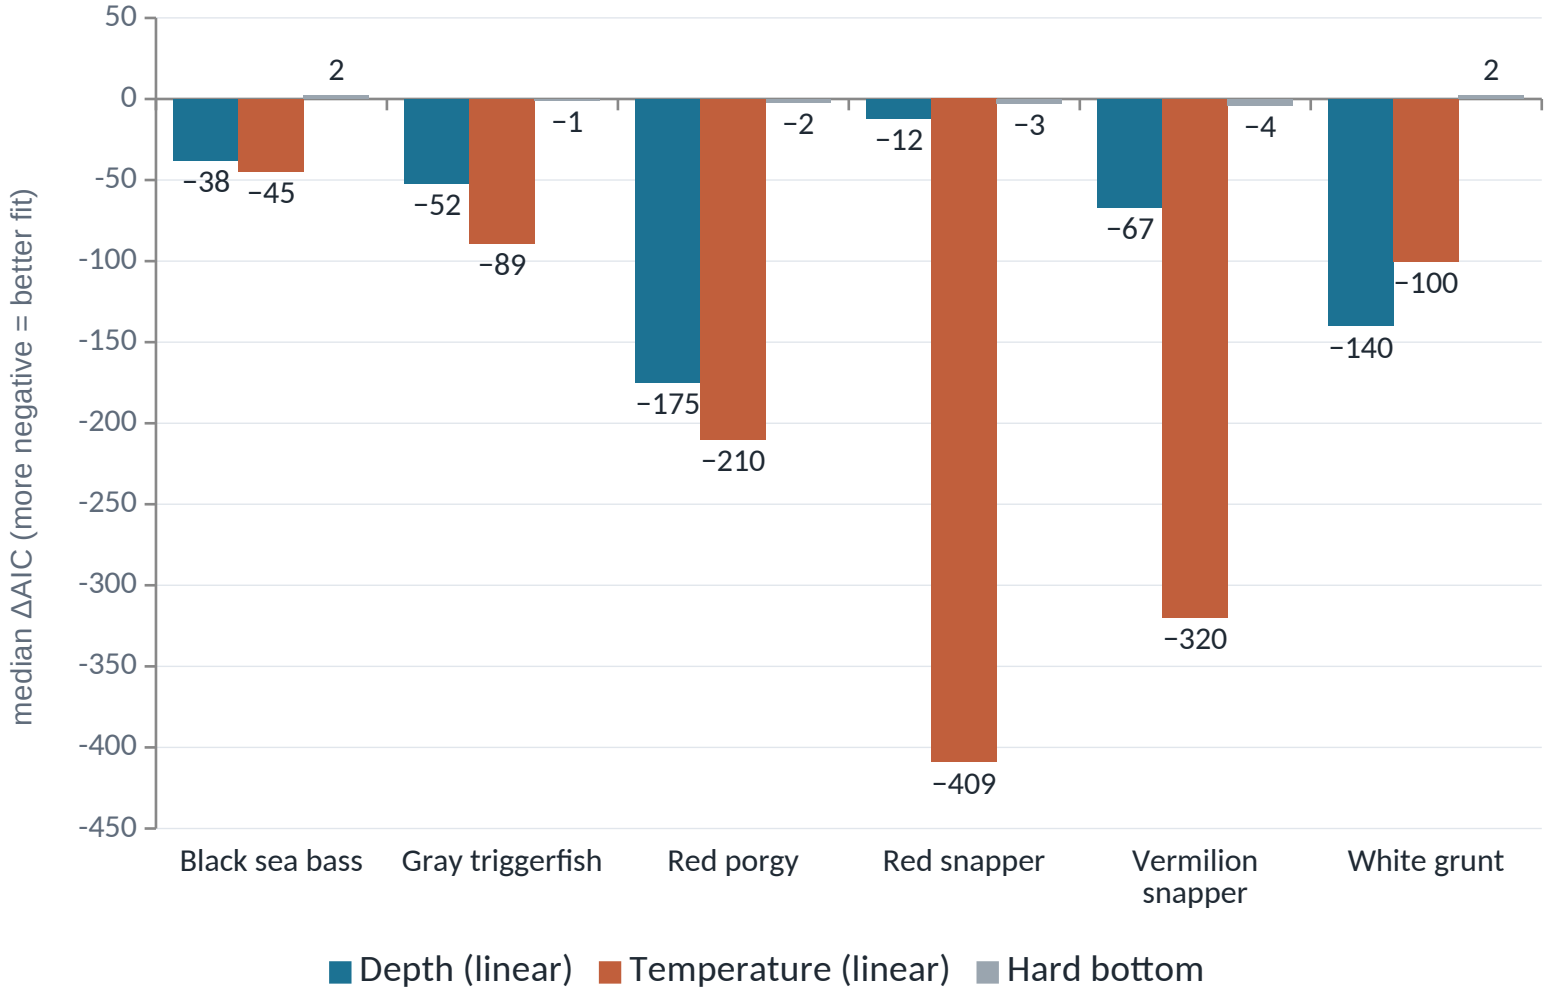
The two criteria do not always agree

Three patterns across the six species

Pattern	Interpretation	Species
AIC improves, field variance falls	Covariate explains spatial structure	Black sea bass, gray triggerfish, red porgy, white grunt
AIC improves, variance flat	Fit gain without attribution of spatial structure	Red snapper (spatial ≈ 0; year-to-year positive)
AIC improves, variance rises	Residual field became more variable despite better fit	Vermilion snapper

Isolating each covariate

Median ΔAIC from adding each covariate, over matched pairs at every level of the other two



Depth and temperature dominate

Temperature reaches −409 AIC units; depth −175.

Hard bottom is negligible

Never more than 4.3 AIC units in either direction, for any species.

Pairs per covariate

Hard bottom 9, depth 6, temperature 6 — ΔAIC and Δ field variance measured per pair.

NC STATE UNIVERSITY

Hard bottom excluded from the final set

Excluded uniformly across the six species rather than conditioning inference on an uneven layer

1 Magnitude

An order of magnitude weaker than depth or temperature in every species — never more than 4.3 AIC units

2 Consistency

Met both criteria in only two of six species; degraded AIC in two others

3 Direction

In two species it improved fit while increasing residual spatial variance — a fit gain without attribution of structure

4 Data quality

Hard-bottom mapping is incomplete across the domain; the weak signal may reflect coverage gaps as much as ecology

Selected model per species

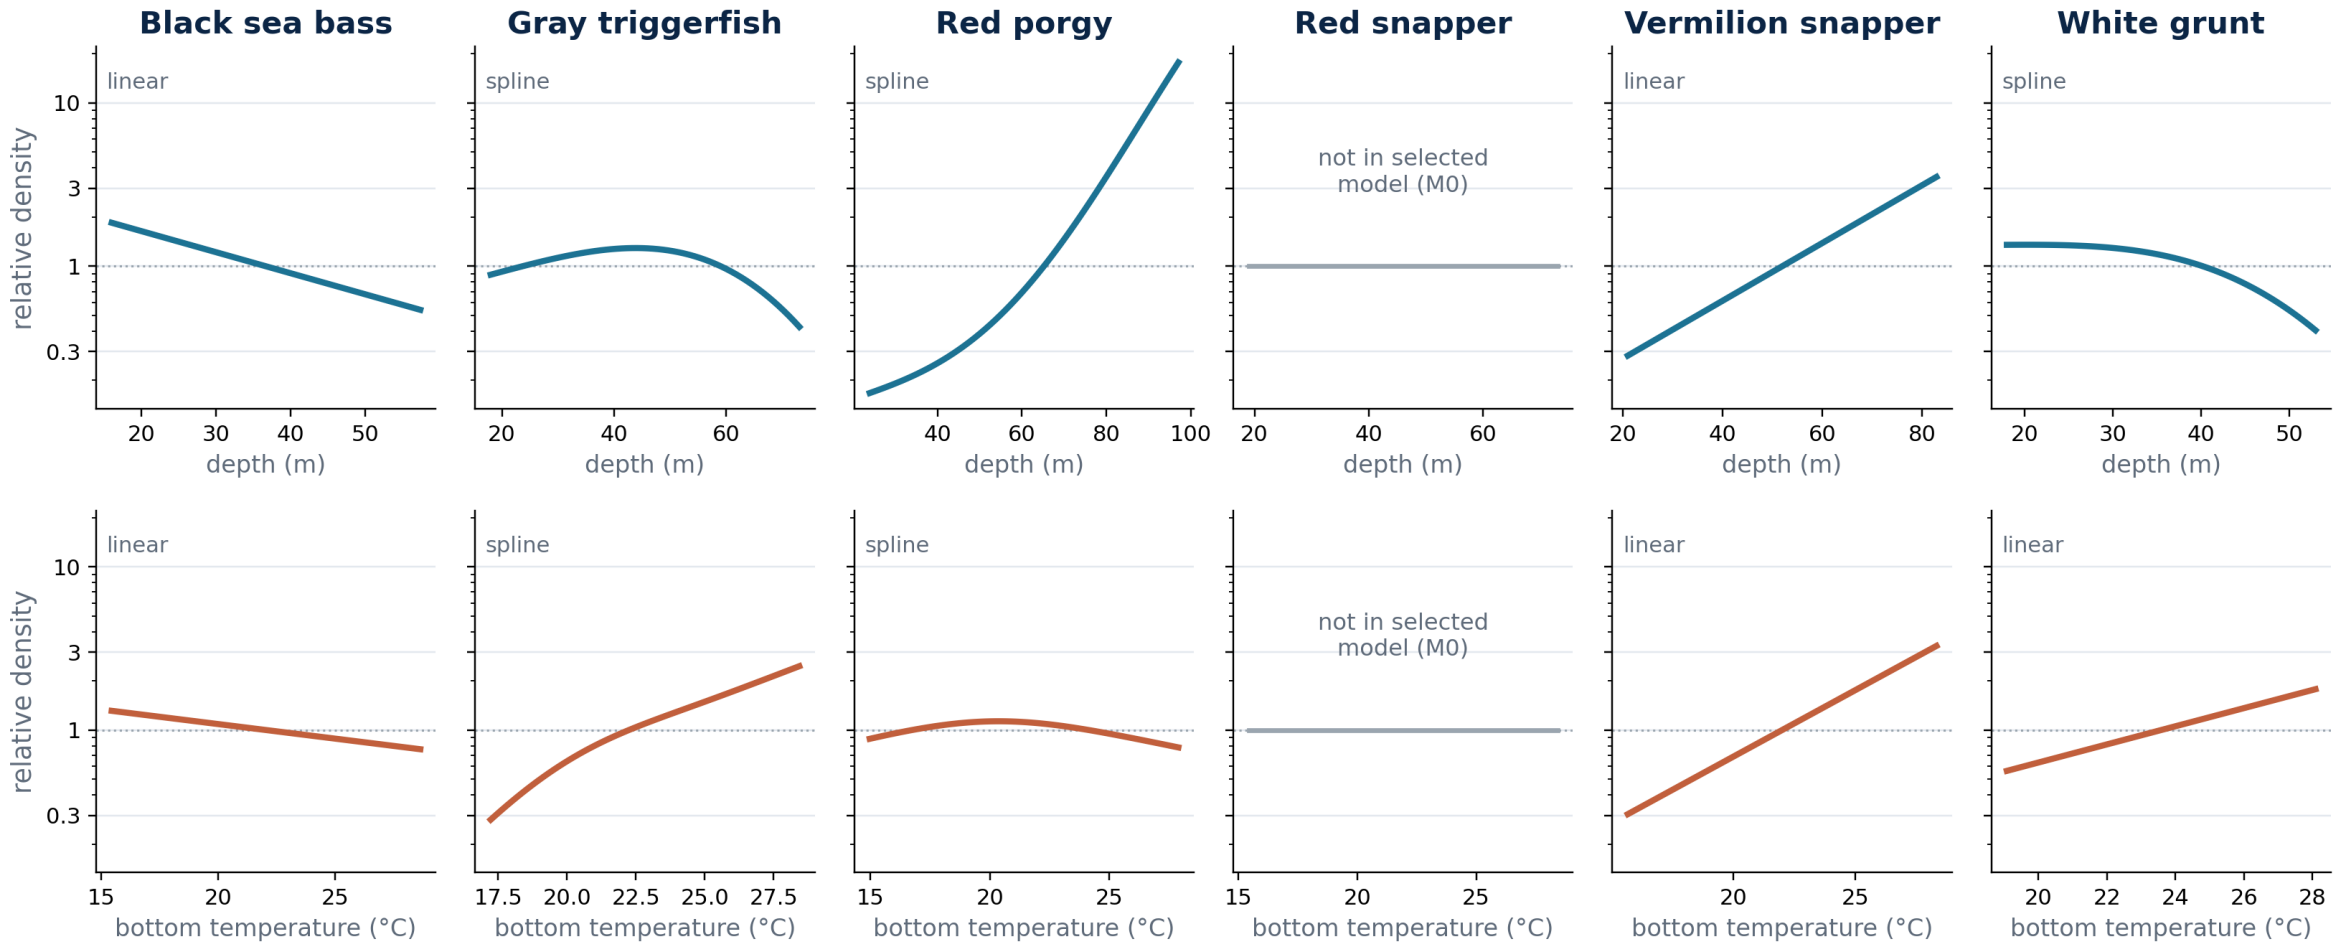
Within the nine-model reduced set: Akaike weight, gap to the next model, and field variance explained

Species	Selected model	Akaike weight	ΔAIC to next	Spatial var. explained	Year-to-year var. explained
Black sea bass	depth + temperature (linear)	0.94	6.9	+14.1 %	−1.1 %
Gray triggerfish	depth + temperature (spline)	1.00	17.6	+9.7 %	+11.5 %
Red porgy	depth + temperature (spline)	0.85	4.2	+16.9 %	+1.5 %
White grunt	depth (spline) + temperature	0.94	5.4	+23.0 %	+2.6 %
Red snapper	temperature (linear)	0.64	1.7	−0.9 %	+1.3 %
Vermilion snapper	depth + temperature (linear)	0.97	8.1	−9.8 %	−3.0 %

Flagged: red snapper (weak, no spatial attribution) and vermilion snapper (residual field more variable). Both are revisited out of sample.

Fitted responses to depth and temperature

Selected model per species; other covariate at its mean, fixed year, location and gear; centred so 1 = average



NC STATE UNIVERSITY

Cross-validation: three designs, three questions

Five folds; models refitted without held-out data; scored by held-out log-likelihood

1 Random folds

- 20 % of events held out at random
- Tests interpolation
- Does the model describe its own data?

2 Spatial folds

- 50-km latitude bands held out
- Tests extrapolation
- Do covariates help where the field cannot interpolate?

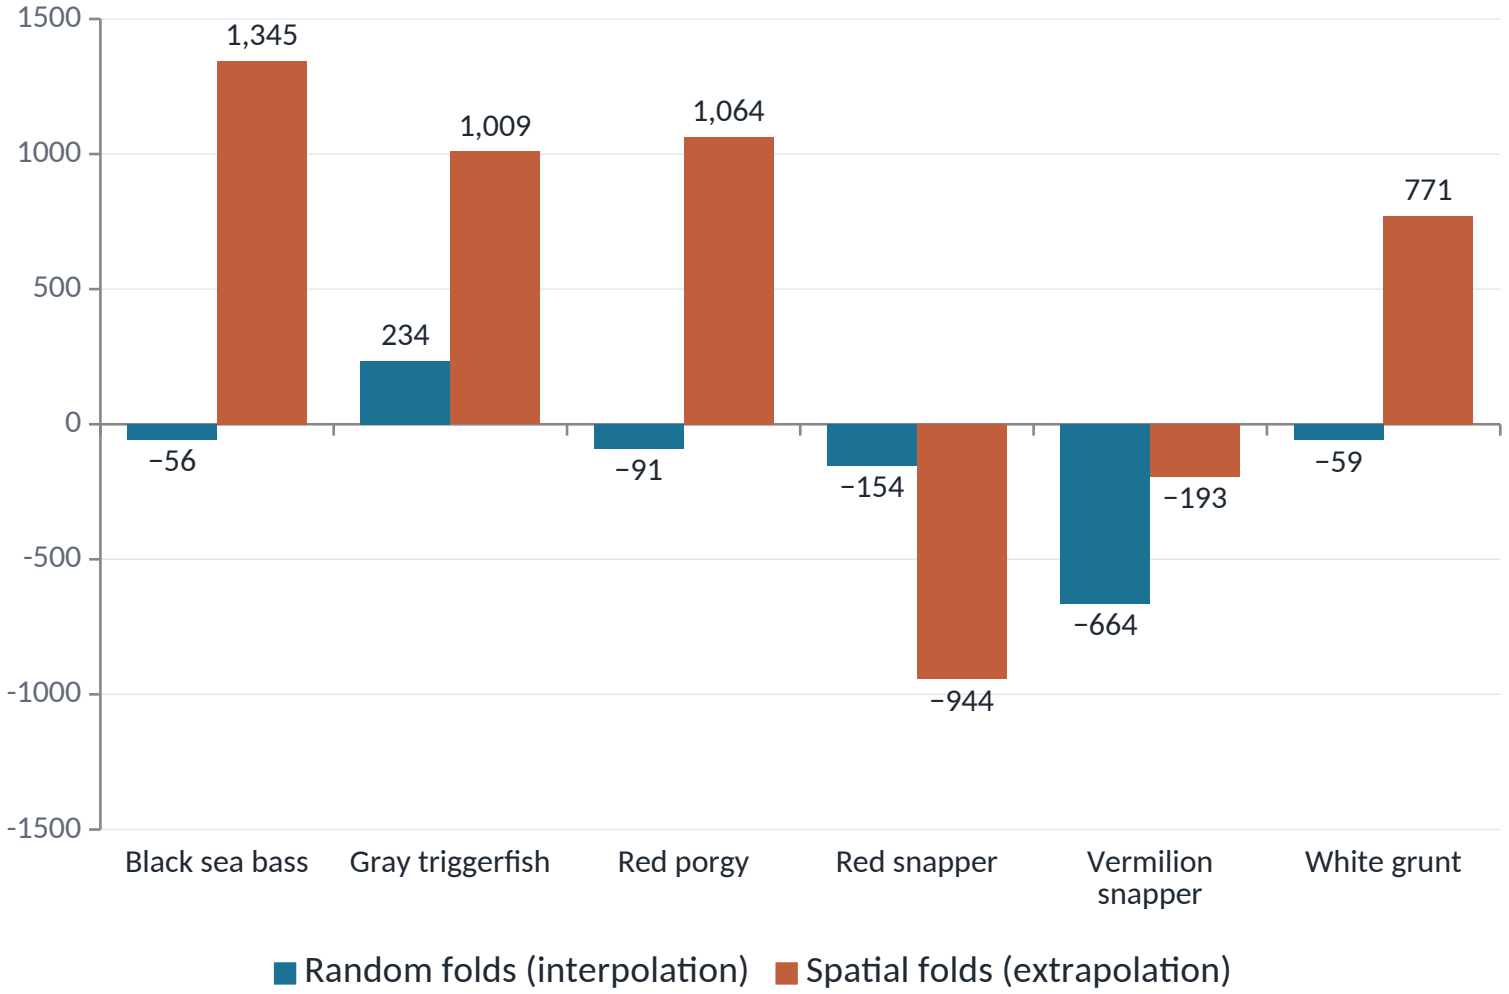
3 Cross-gear folds

- One gear held out; other kept
- Tests transfer between gears
- Does sharing let video inform trap, and vice versa?

In every design: null model, AIC-selected model, and a gear-independent model with the same covariates. Spatial folds scored by log predictive density over 200 field draws.

CV result 1 — do the covariates help?

Random and spatial folds. Held-out log-likelihood gain of the selected covariate model over the null, summed over 43,336 observations



Interpolation: no gain

The field already carries what depth and temperature add.

Extrapolation: real gain

+770 to +1,350 for four species: the model transports better.

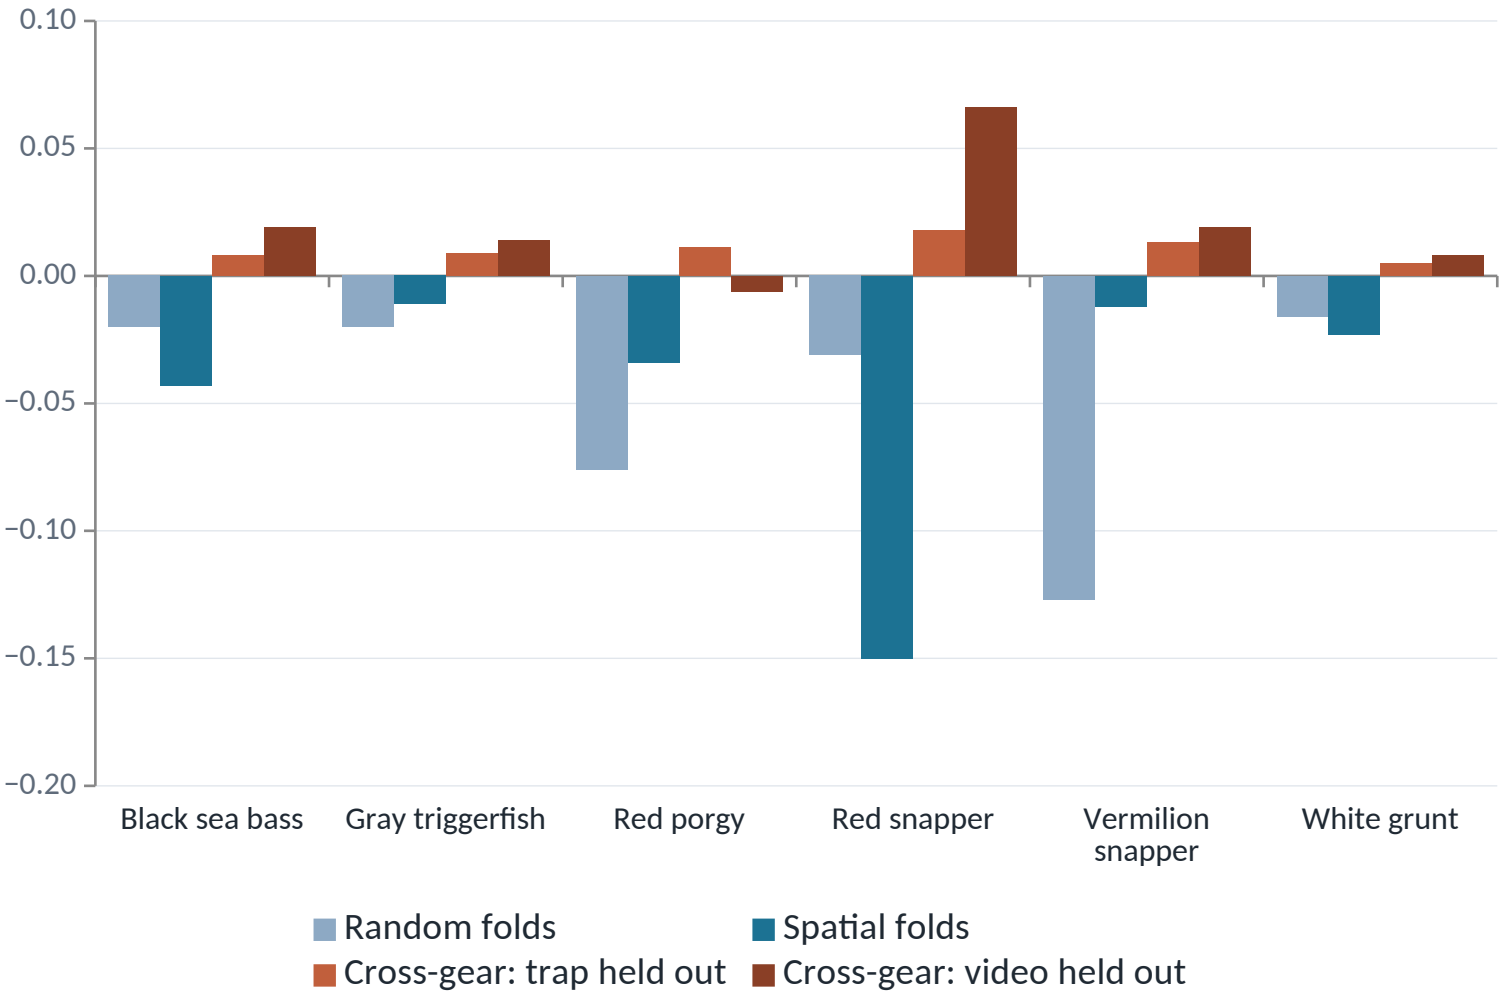
Two exceptions

Vermilion snapper: slightly hurt. Red snapper: worse than null → null model used.

NC STATE UNIVERSITY

CV result 2 — does sharing fields help?

All three designs. Shared-field minus gear-independent score per observation; positive favours sharing



Within a gear

Independent fields fit slightly better: each gear has its own structure.

Across gears

Shared fields predict the missing gear better: 6 of 6 (trap out), 5 of 6 (video out).

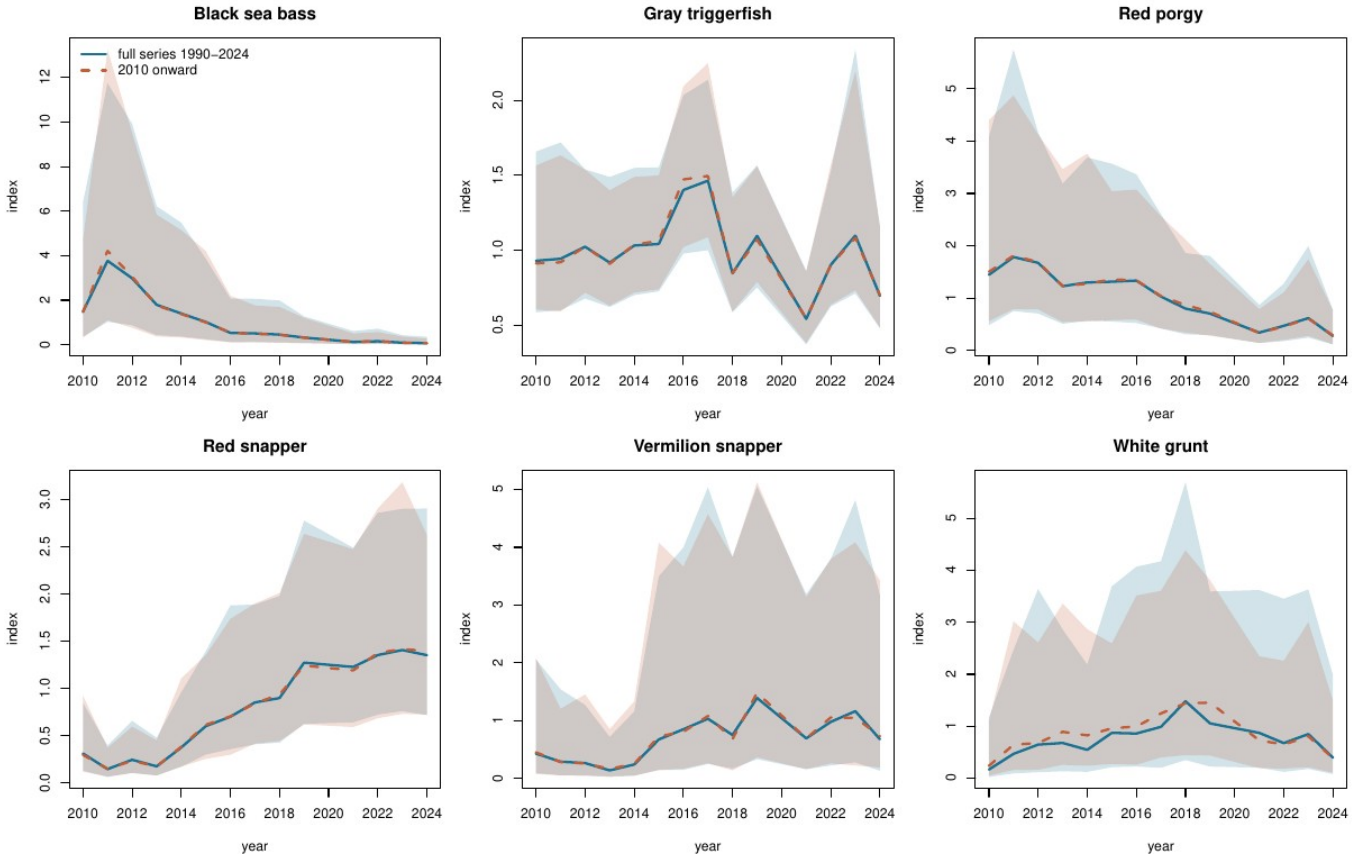
Why that settles it

1990–2009 has trap only. That is the cross-gear case.

NC STATE UNIVERSITY

Sensitivity 1 — refit to 2010–2024 only

Each species refitted to 2010–2024 only and compared with the full fit over the same years



Index 2010–2024: full-series fit (solid) vs 2010-onward fit (dashed), with 95 % bands.

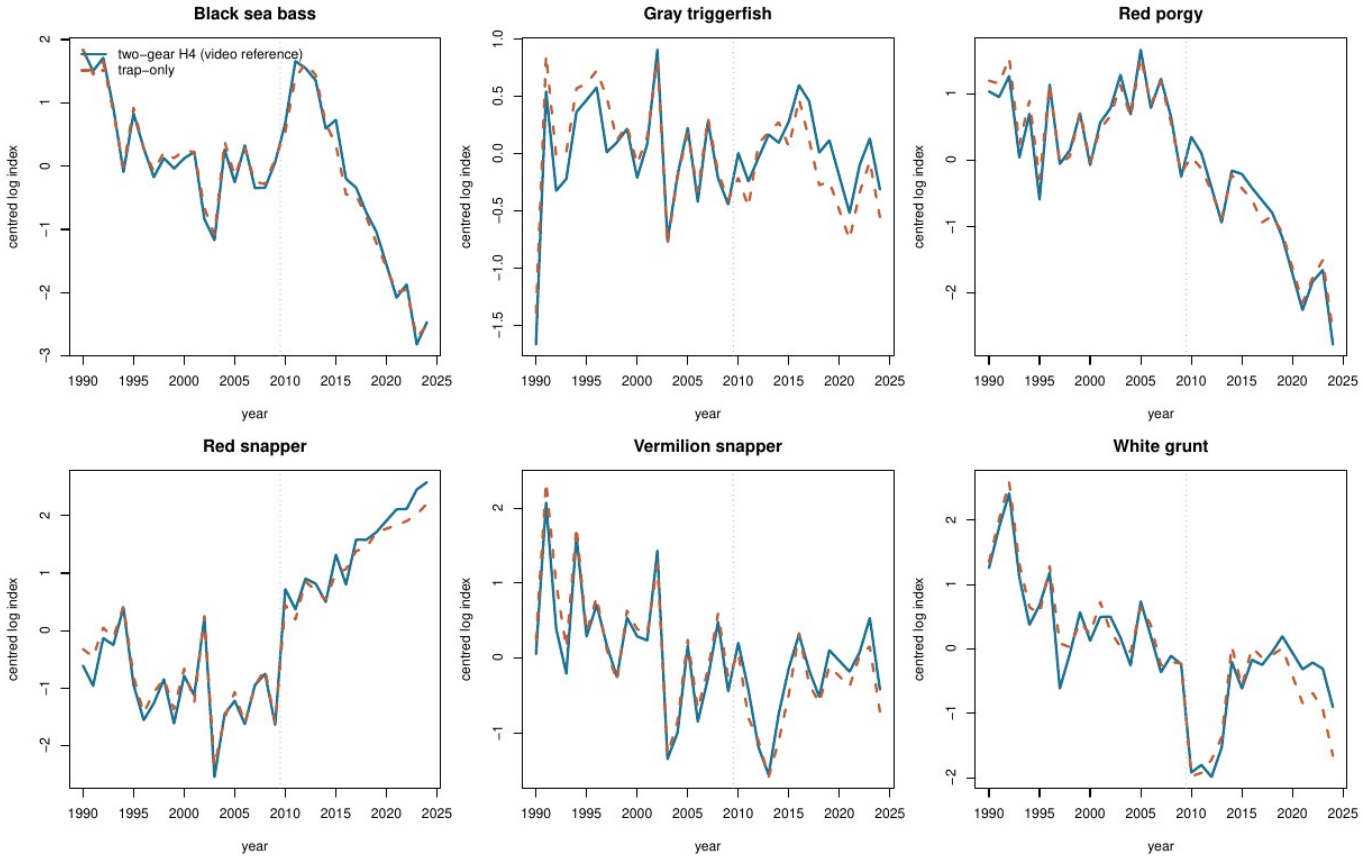
Species	Index r	COG r	Temp r
Black sea bass	0.998	0.988	0.997
Gray triggerfish	0.996	0.998	1.000
Red porgy	0.999	0.978	0.993
Red snapper	0.999	0.994	1.000
Vermilion snapper	0.989	0.976	0.997
White grunt	0.903	0.881	0.982

Reading

- Era-only medians lie inside the full-series intervals in every year
- Early data add years, not bias

Sensitivity 2 — what does the video change?

Each species refitted to trap data alone and compared with the two-gear model



Centred log index at sampled sites: two-gear model (solid) vs trap-only (dashed); dotted line = 2010.

Species	Map r	Index r	Amplitude
Black sea bass	0.999	0.99	0.98
Gray triggerfish	0.94	0.91	0.94
Red porgy	0.96	0.99	0.97
Red snapper	0.92	0.99	0.90
Vermilion snapper	0.82	0.97	1.04
White grunt	0.88	0.96	1.03

Reading

- For these six, the trap alone gives nearly the same index
- Video's value: gear separation, cross-gear transfer, and species traps barely catch

NC STATE UNIVERSITY

The six production models

H4 structure, 5-km mesh, negative binomial, fixed year effects

Species	Covariates	Basis	Range (km)	Caveat
Black sea bass	depth L, temp L	AIC + CV	20.5	
Gray triggerfish	depth S, temp S	AIC + CV	16.8	
Red porgy	depth S, temp S	AIC + CV	23.0	
Red snapper	none	**CV overruled AIC**	24.2	covariate model worse than null out of sample
Vermilion snapper	depth L, temp L	AIC	18.2	covariates weakly supported
White grunt	depth S, temp L	AIC + CV	41.5	widest field, noisiest index

Uncertainty: 1,000 draws of the random effects with fixed effects at their estimates. Trap loadings on the annual field 0.56–0.87: the trap responds more weakly to abundance change than the video.

Products: the indicator set

Every indicator comes from 1,000 draws of the fitted fields, with a 95 % interval; trends also report P(trend > 0)

Abundance

Annual relative abundance index, 1990–2024, on the population scale (gear effects removed)

Location

Centre of gravity, northing and easting; year-to-year displacement; northern and southern range edges (5th and 95th percentiles)

Extent

Effective area occupied, the area that would hold the population at its mean density

Habitat occupied

Density-weighted mean depth and bottom temperature; climate tracking, whether the species follows isotherms or stays put

Trends

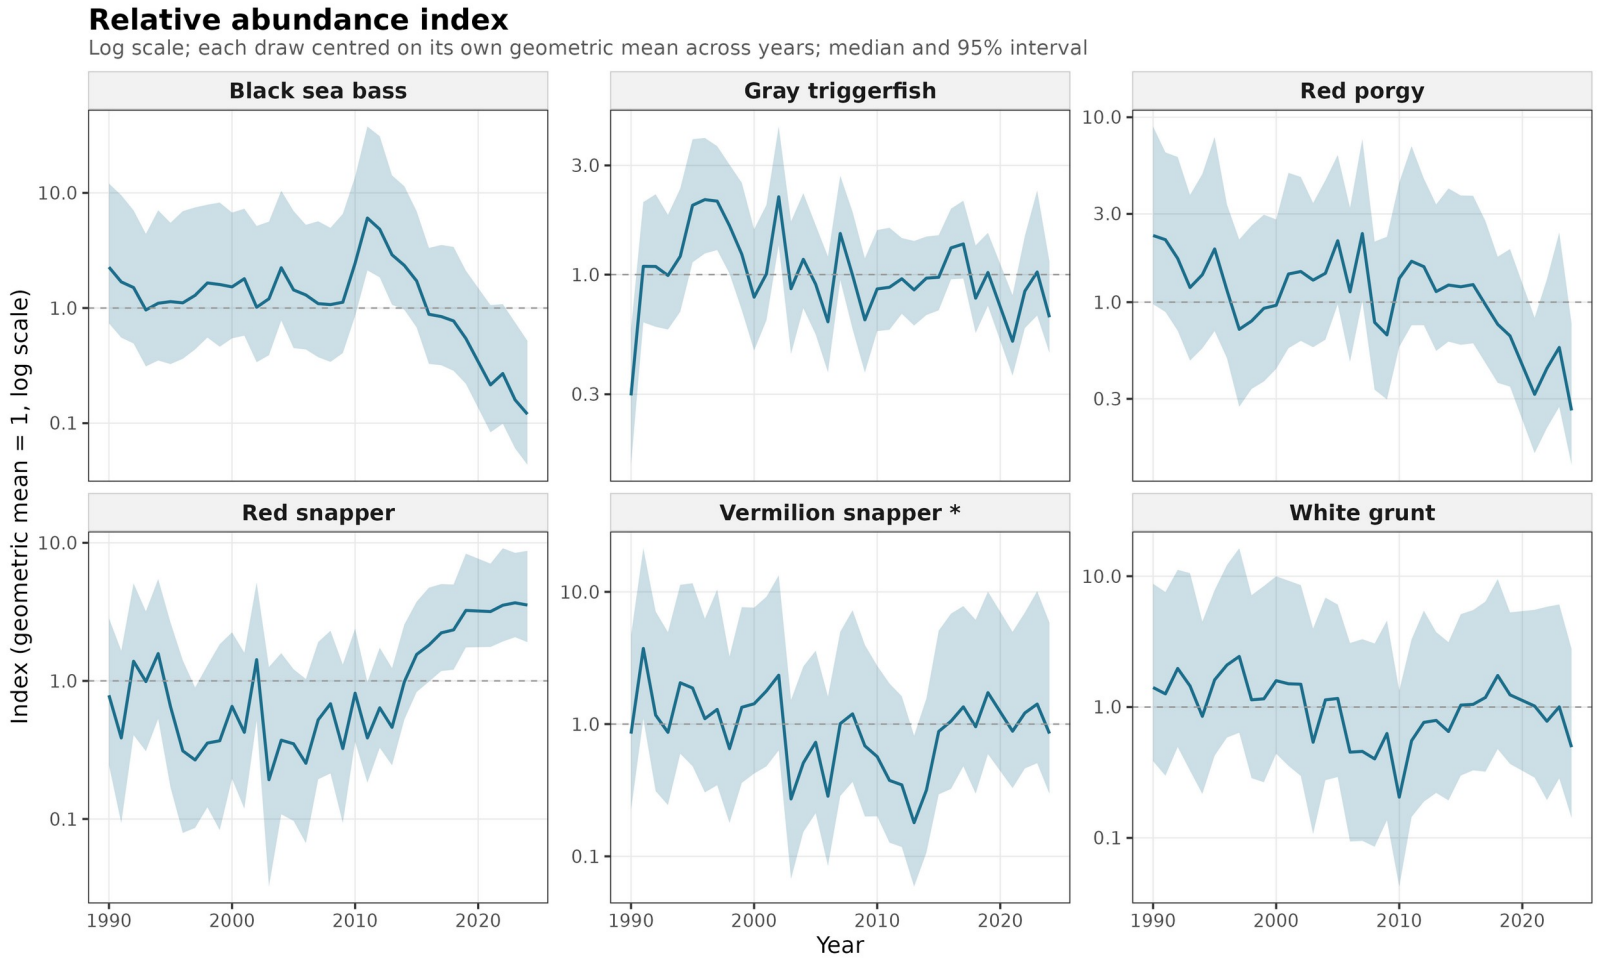
Linear trend rate of each indicator with P(trend > 0); terminal change, the last five years against the first five

Habitat (EFH)

Persistent density, abundance tiers, stability and depth–temperature envelopes

Products: relative abundance indices, 1990–2024

Log scale, centred per draw; median and 95 % interval from 1,000 draws



Red snapper

Rising since ~2013 to the highest levels of the series

Black sea bass

Peak 2011–12, then decline to a series low in 2024

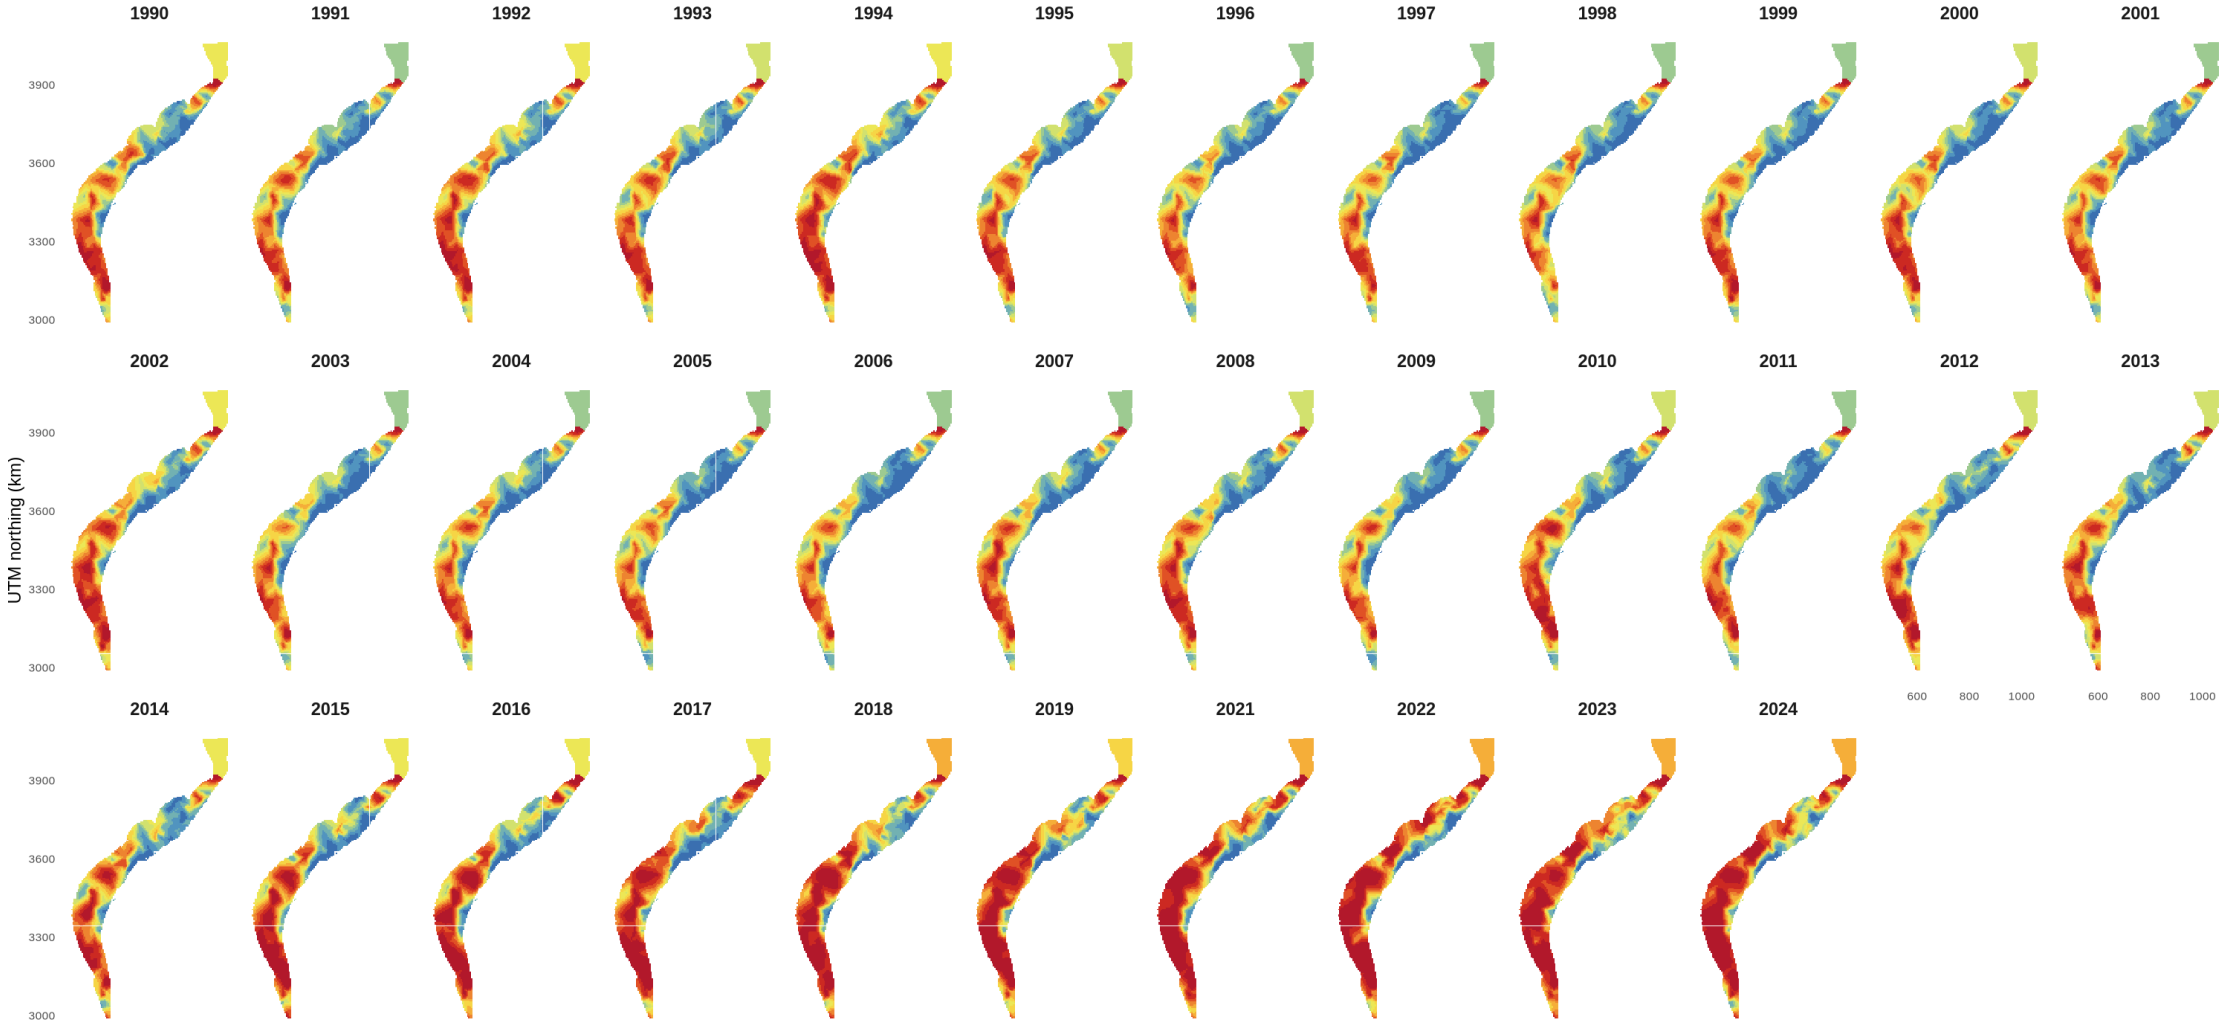
The others

Flat within wide intervals; pre-2010 values rest on trap data alone

NC STATE UNIVERSITY

Products: density surfaces for every survey year

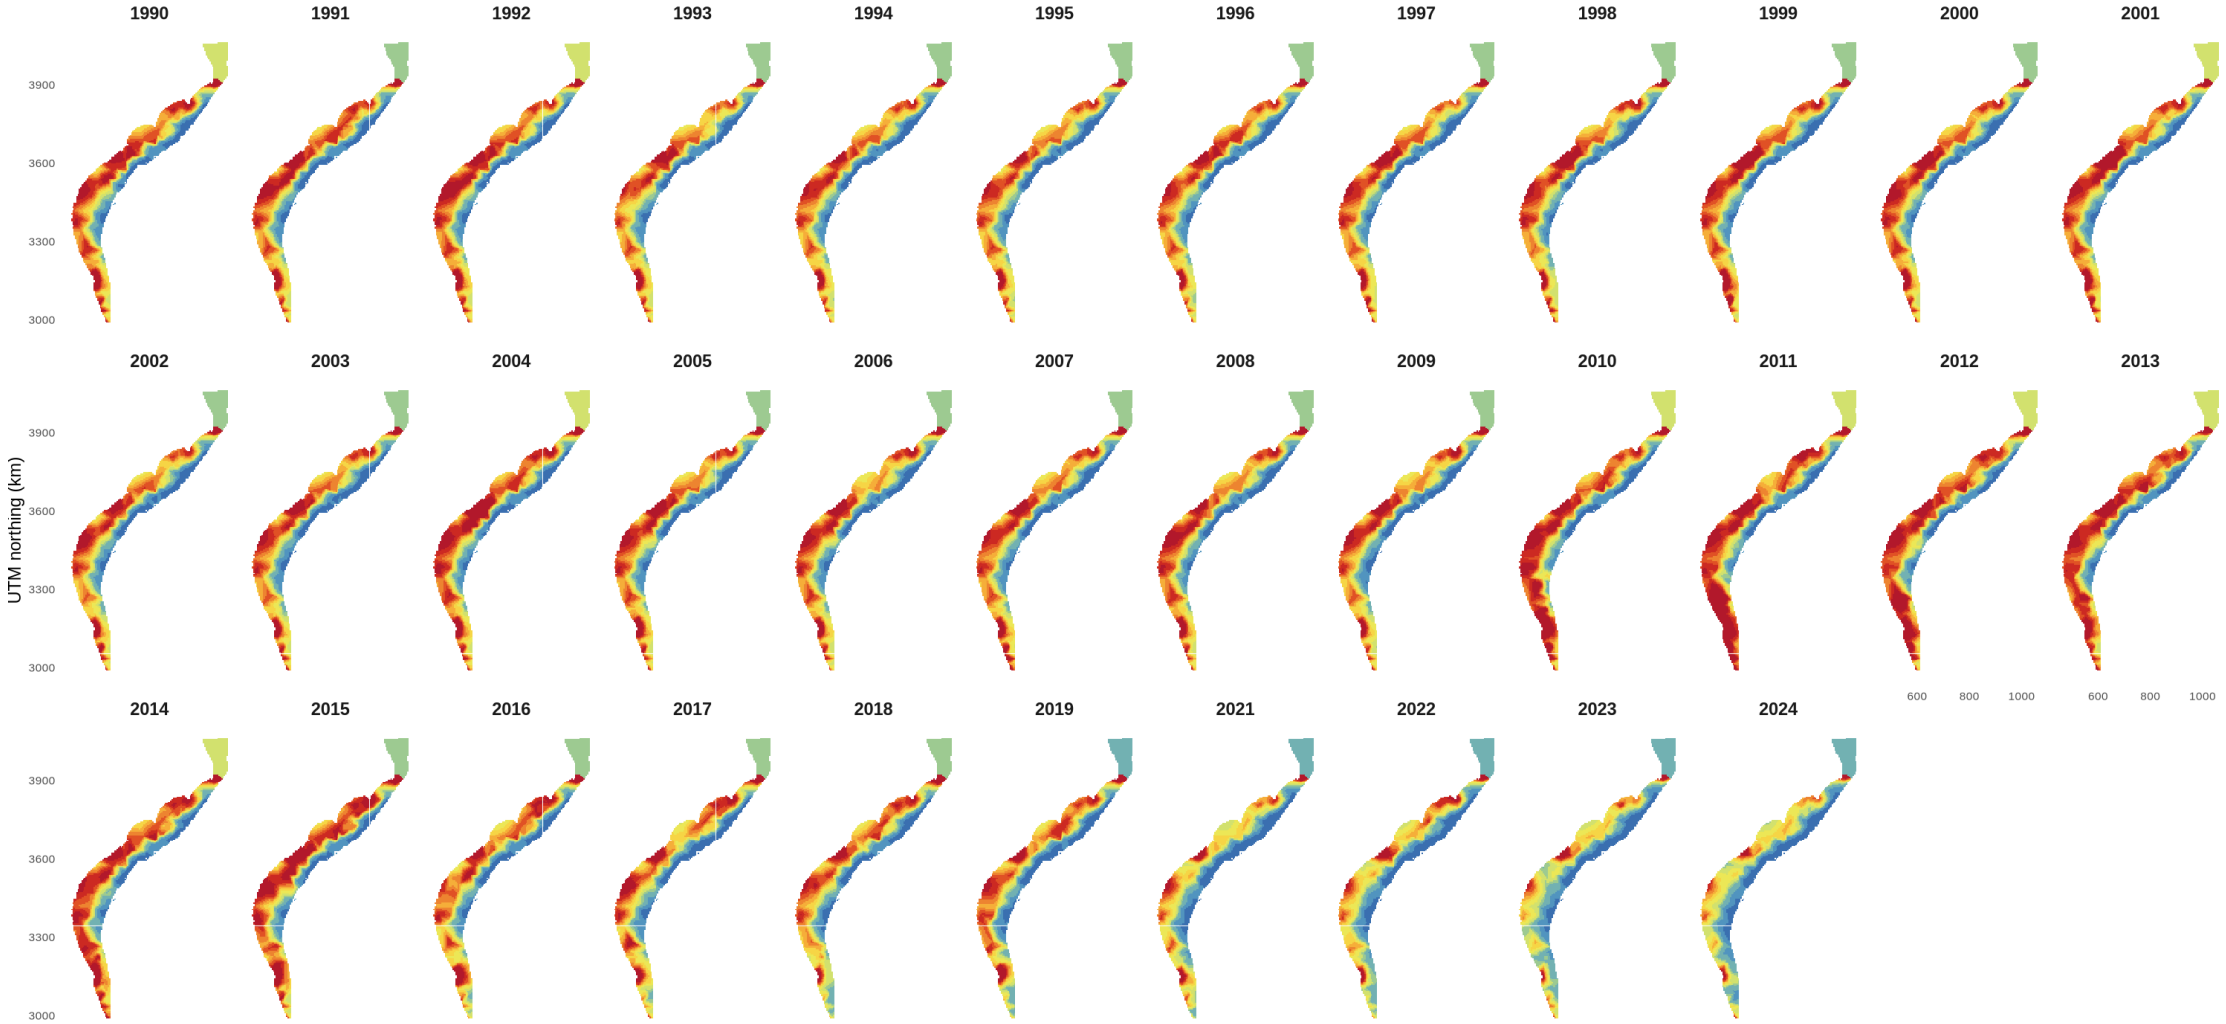
Red snapper, 5-km cells; common quantile colour classes across years (blue = low, red = high)



NC STATE UNIVERSITY

Products: density surfaces for every survey year

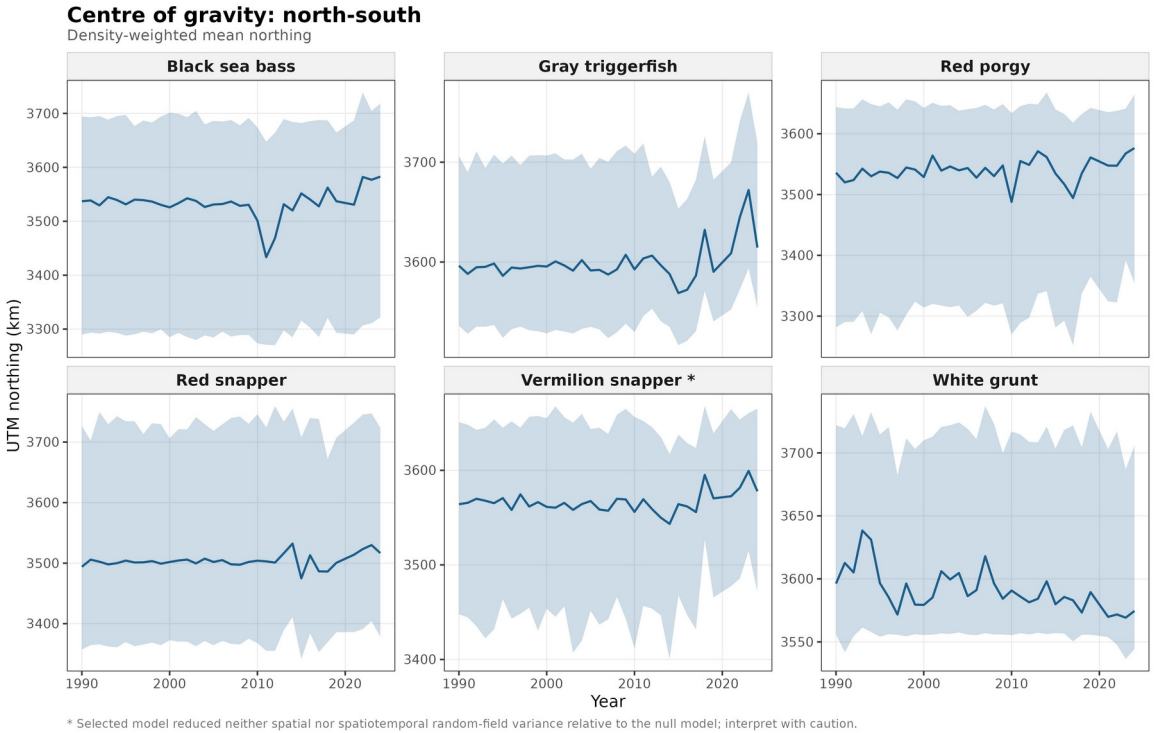
Black sea bass, 5-km cells; common quantile colour classes across years (blue = low, red = high)



NC STATE UNIVERSITY

Products: distribution indicators and trends

Centre of gravity (north–south) with 95 % intervals; linear trends 1990–2024 from per-draw slopes



Species	North (km/yr)	Temp (°C/yr)	P(warm)
Black sea bass	0.6 (−2.0, 3.5)	0.018 (0.010, 0.026)	1.00
Gray triggerfish	1.0 (−0.2, 2.2)	0.022 (0.017, 0.026)	1.00
Red porgy	0.7 (−1.2, 3.1)	0.038 (0.008, 0.076)	0.99
Red snapper	0.5 (−1.5, 2.6)	0.024 (0.014, 0.038)	1.00
Vermilion snapper	0.5 (−0.9, 2.2)	0.031 (0.017, 0.049)	1.00
White grunt	−0.8 (−2.1, 0.3)	0.029 (0.015, 0.037)	1.00

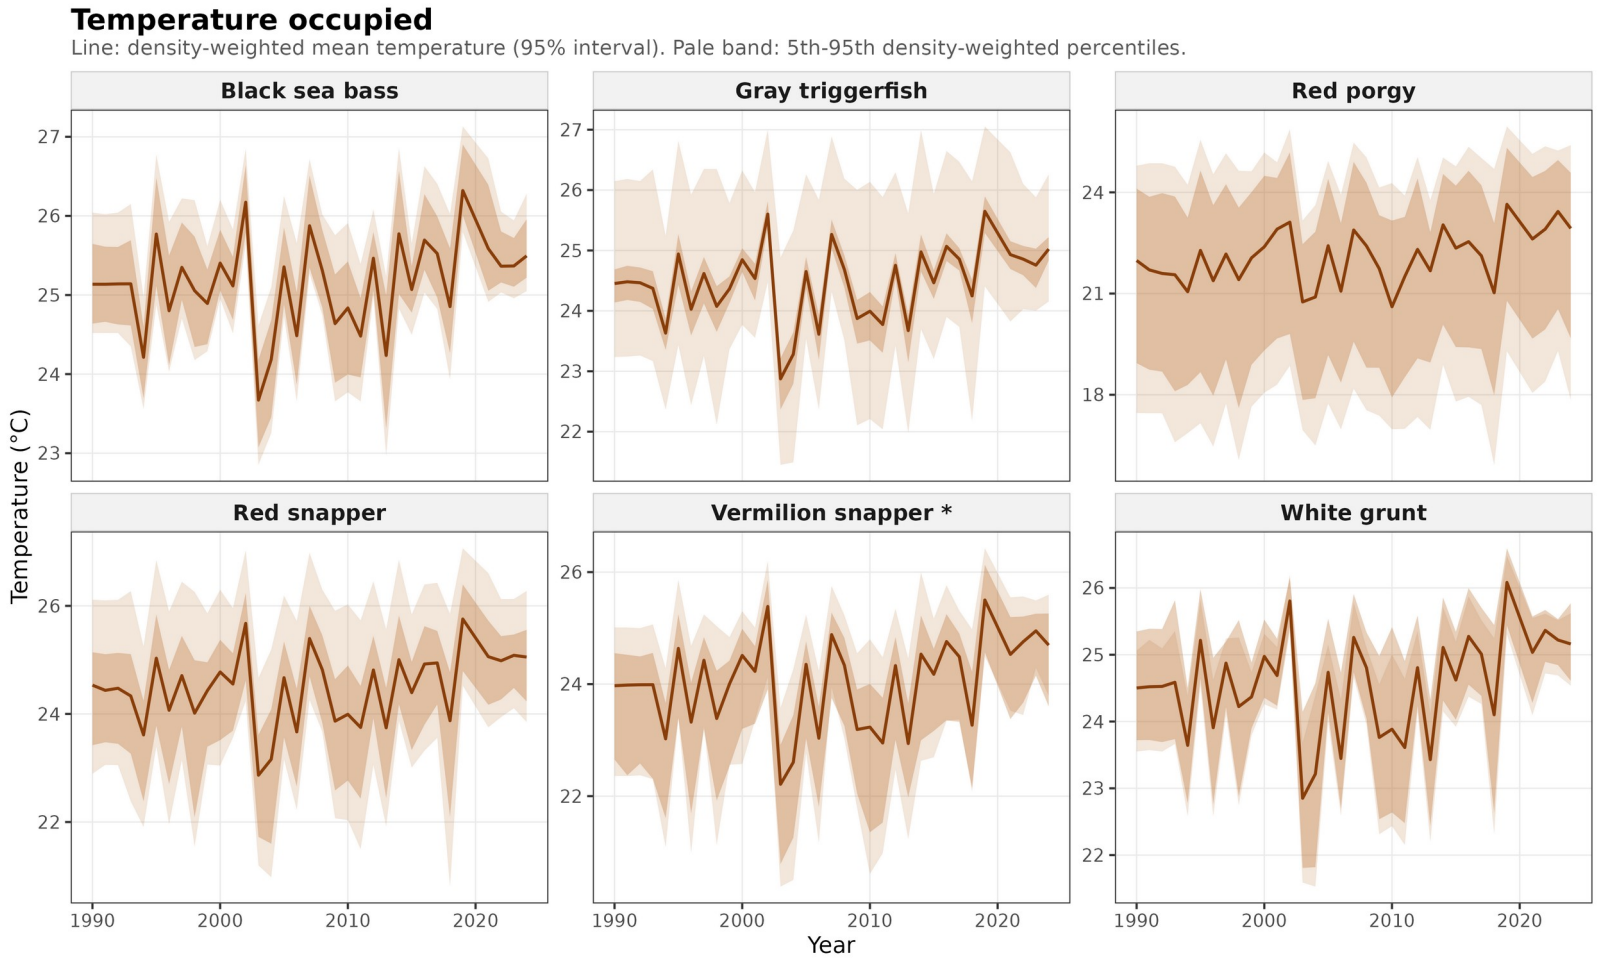
Median (95 % interval) of per-draw slopes; P = probability of a positive trend.

No detectable northward or depth shift in any species. Occupied bottom temperature is rising in all six — the community is warming in place.

NC STATE UNIVERSITY

Every species occupies warmer water than in 1990

Density-weighted mean bottom temperature of occupied habitat, 1990–2024



Temperature is treated as known; no uncertainty is propagated from the environmental fields.

* Selected model reduced neither spatial nor spatiotemporal random-field variance relative to the null model; interpret with caution.

6 / 6

species warming
(all P ≥ 0.99)

+0.02–0.04

°C per year:
0.6–1.3 °C over the series

NC STATE UNIVERSITY

Essential fish habitat: the four products

1 Persistent density

The across-year average of the annual surfaces: where the species is found year after year, net of year-to-year shifts. The base layer for everything below.

2 Abundance tiers

Cells ranked by persistent density and accumulated from the top. The smallest set holding 95 % of abundance is the habitat area; 50 % is the core; 25 % marks hot spots.

3 Stability

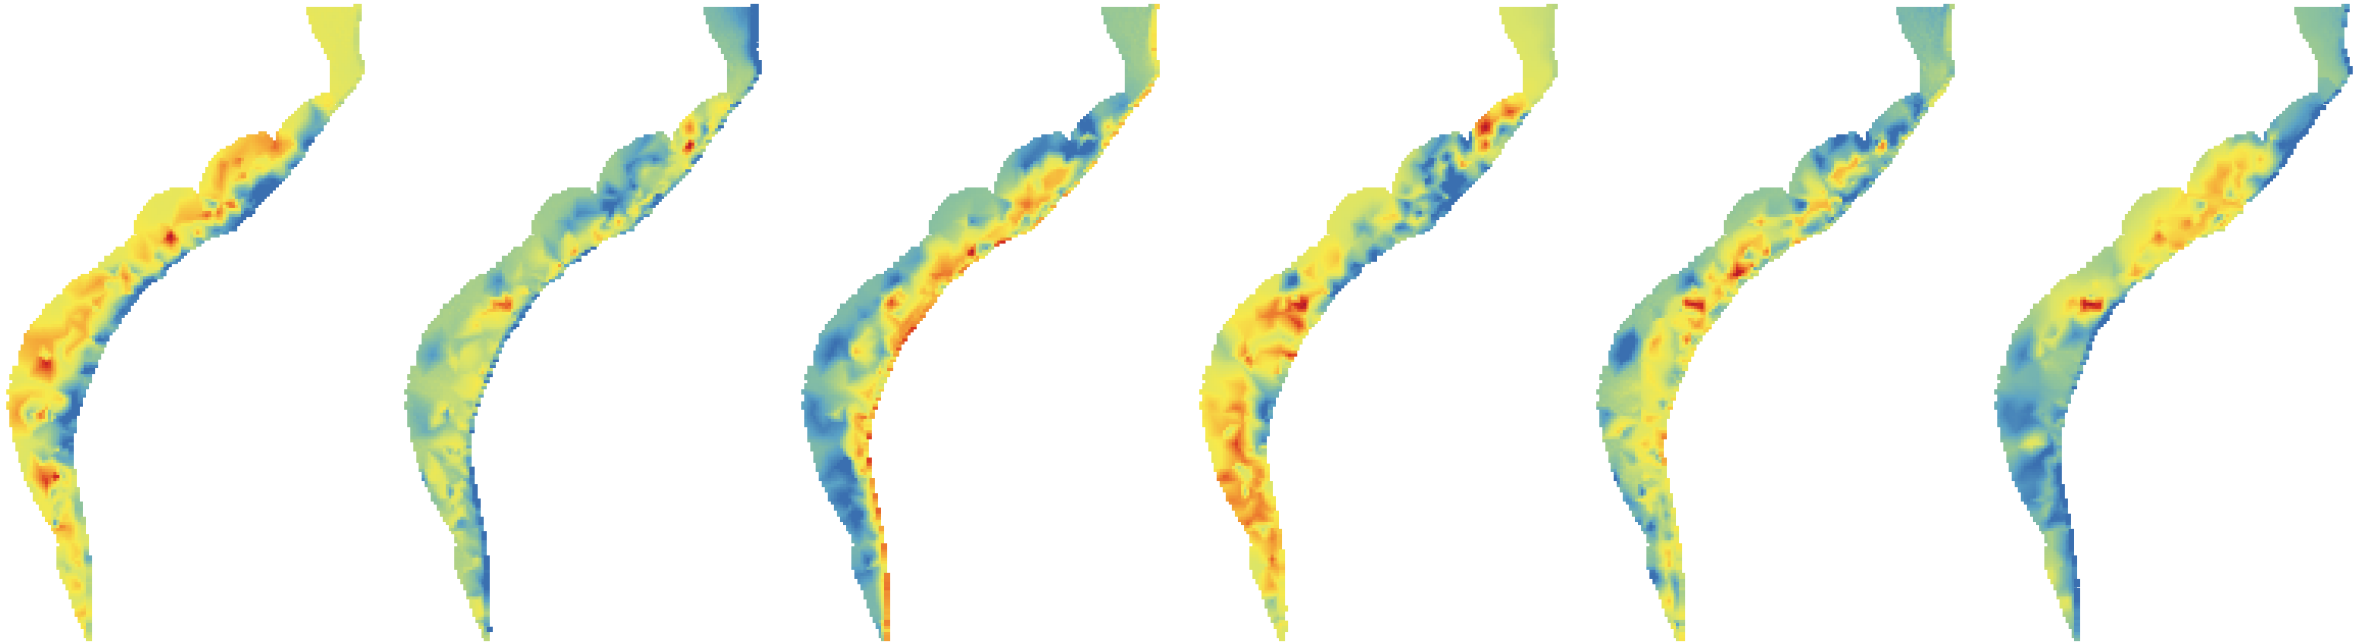
Tiers recomputed separately for each year; for each cell, the share of years it falls inside that year's 95 % tier. Separates habitat that is always used from habitat that drifts.

4 Depth–temperature envelope

Density-weighted ranges of depth and bottom temperature holding 95 % and 50 % of persistent abundance: the text-based statement a designation needs.

Persistent density for the six species

Across-year average of the annual single-species (H4) surfaces, 1990–2024; log scale, each species from low to its peak



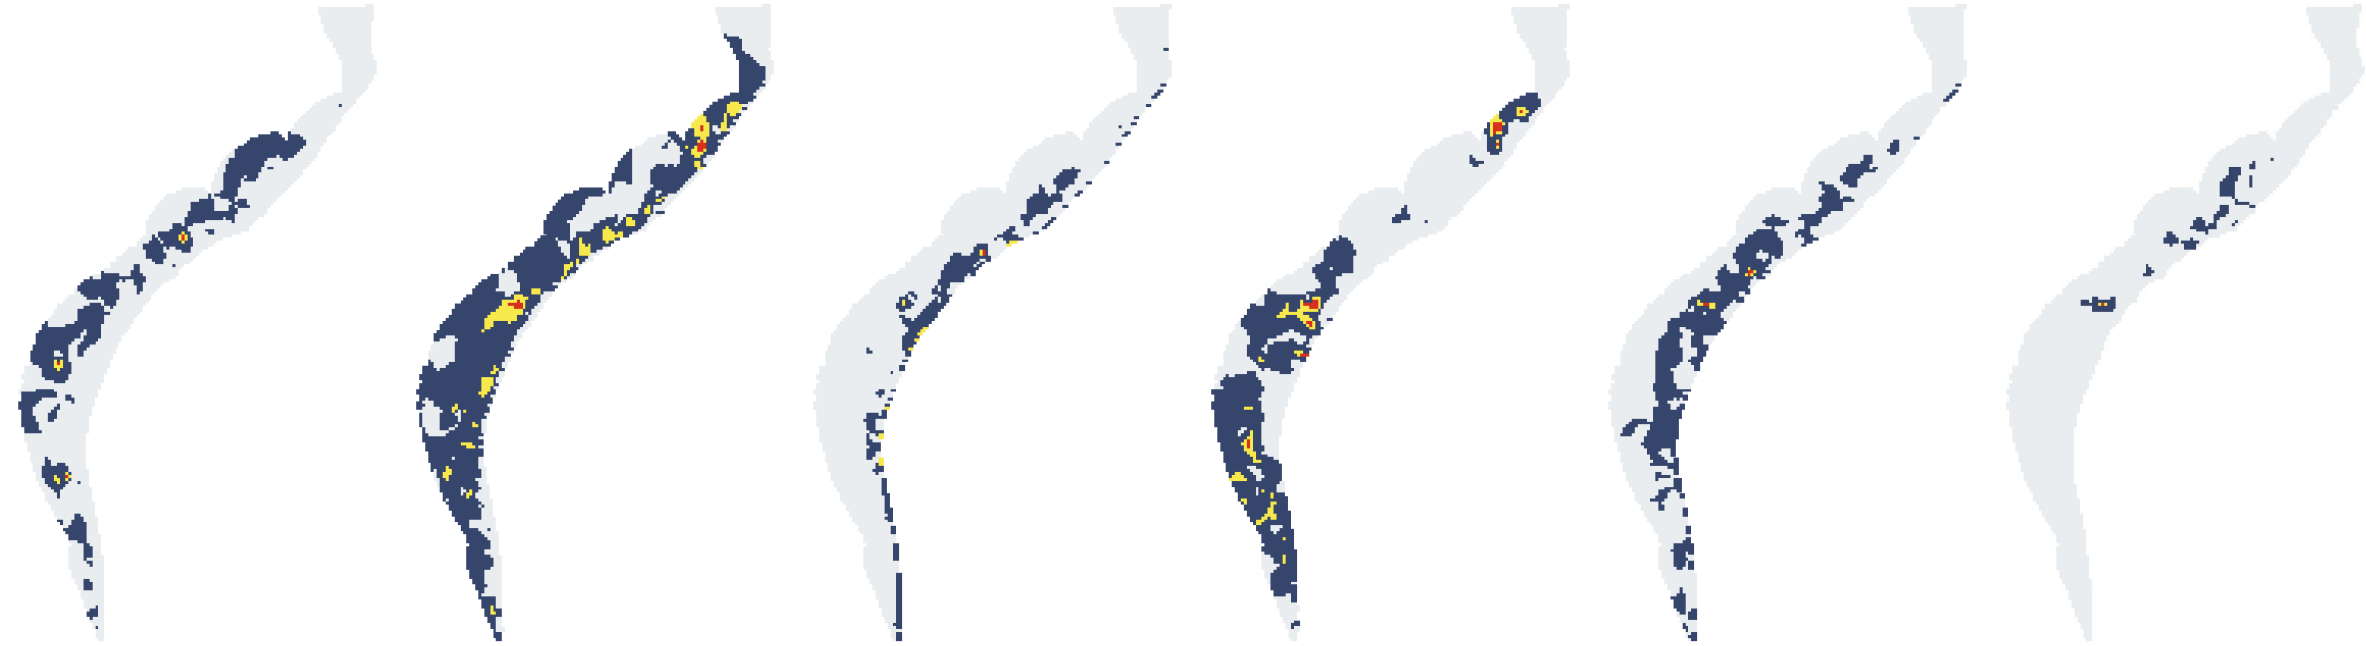
Black sea bass | **Gray triggerfish** | **Red porgy** | **Red snapper** | **Vermilion snapper** | **White grunt**



NC STATE UNIVERSITY

Habitat tiers for the six species

Smallest area holding 95 % (habitat), 50 % (core) and 25 % (hot spots) of persistent predicted abundance, 1990–2024



Black sea bass
95 % area = 31 % of domain

Gray triggerfish
95 % area = 71 % of domain

Red porgy
95 % area = 13 % of domain

Red snapper
95 % area = 44 % of domain

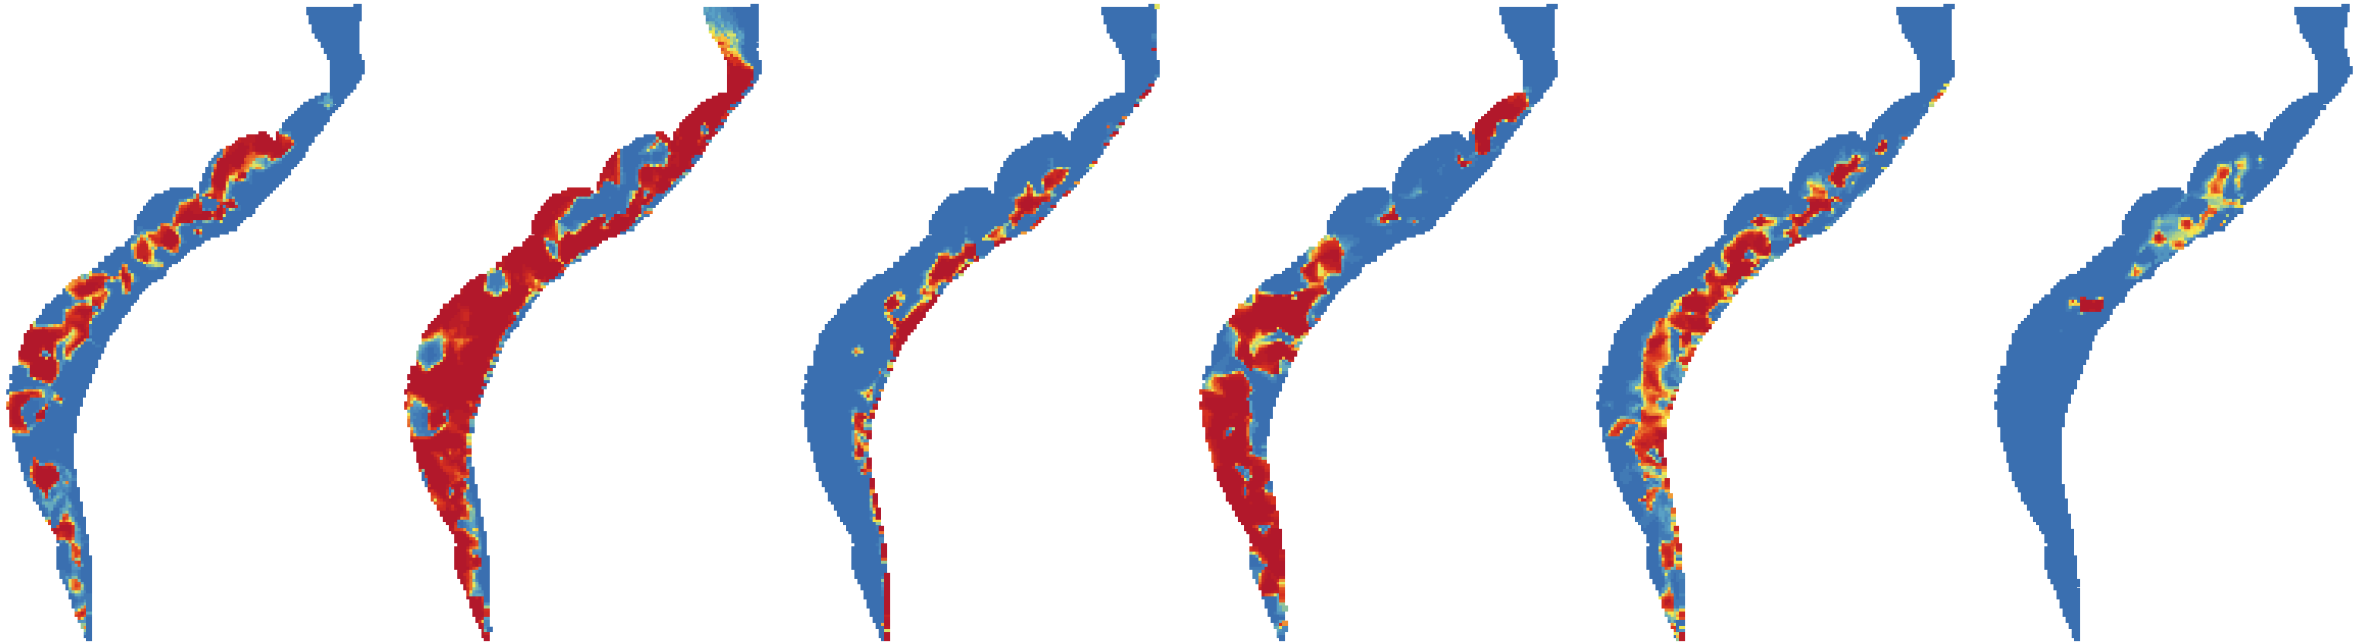
Vermilion snapper
95 % area = 30 % of domain

White grunt
95 % area = 5 % of domain

hot spot (25 %)
core (50 %)
habitat (95 %)
outside

Stability of the 95 % tier for the six species

Share of the 34 survey years in which each cell falls inside that year's 95 % habitat tier



Black sea bass
mean inside tier = 0.9

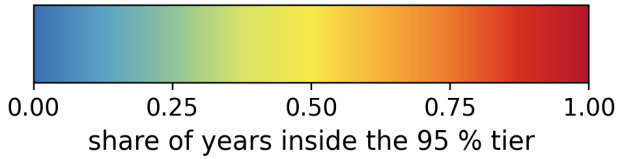
Gray triggerfish
mean inside tier = 1.0

Red porgy
mean inside tier = 0.9

Red snapper
mean inside tier = 0.9

Vermilion snapper
mean inside tier = 0.8

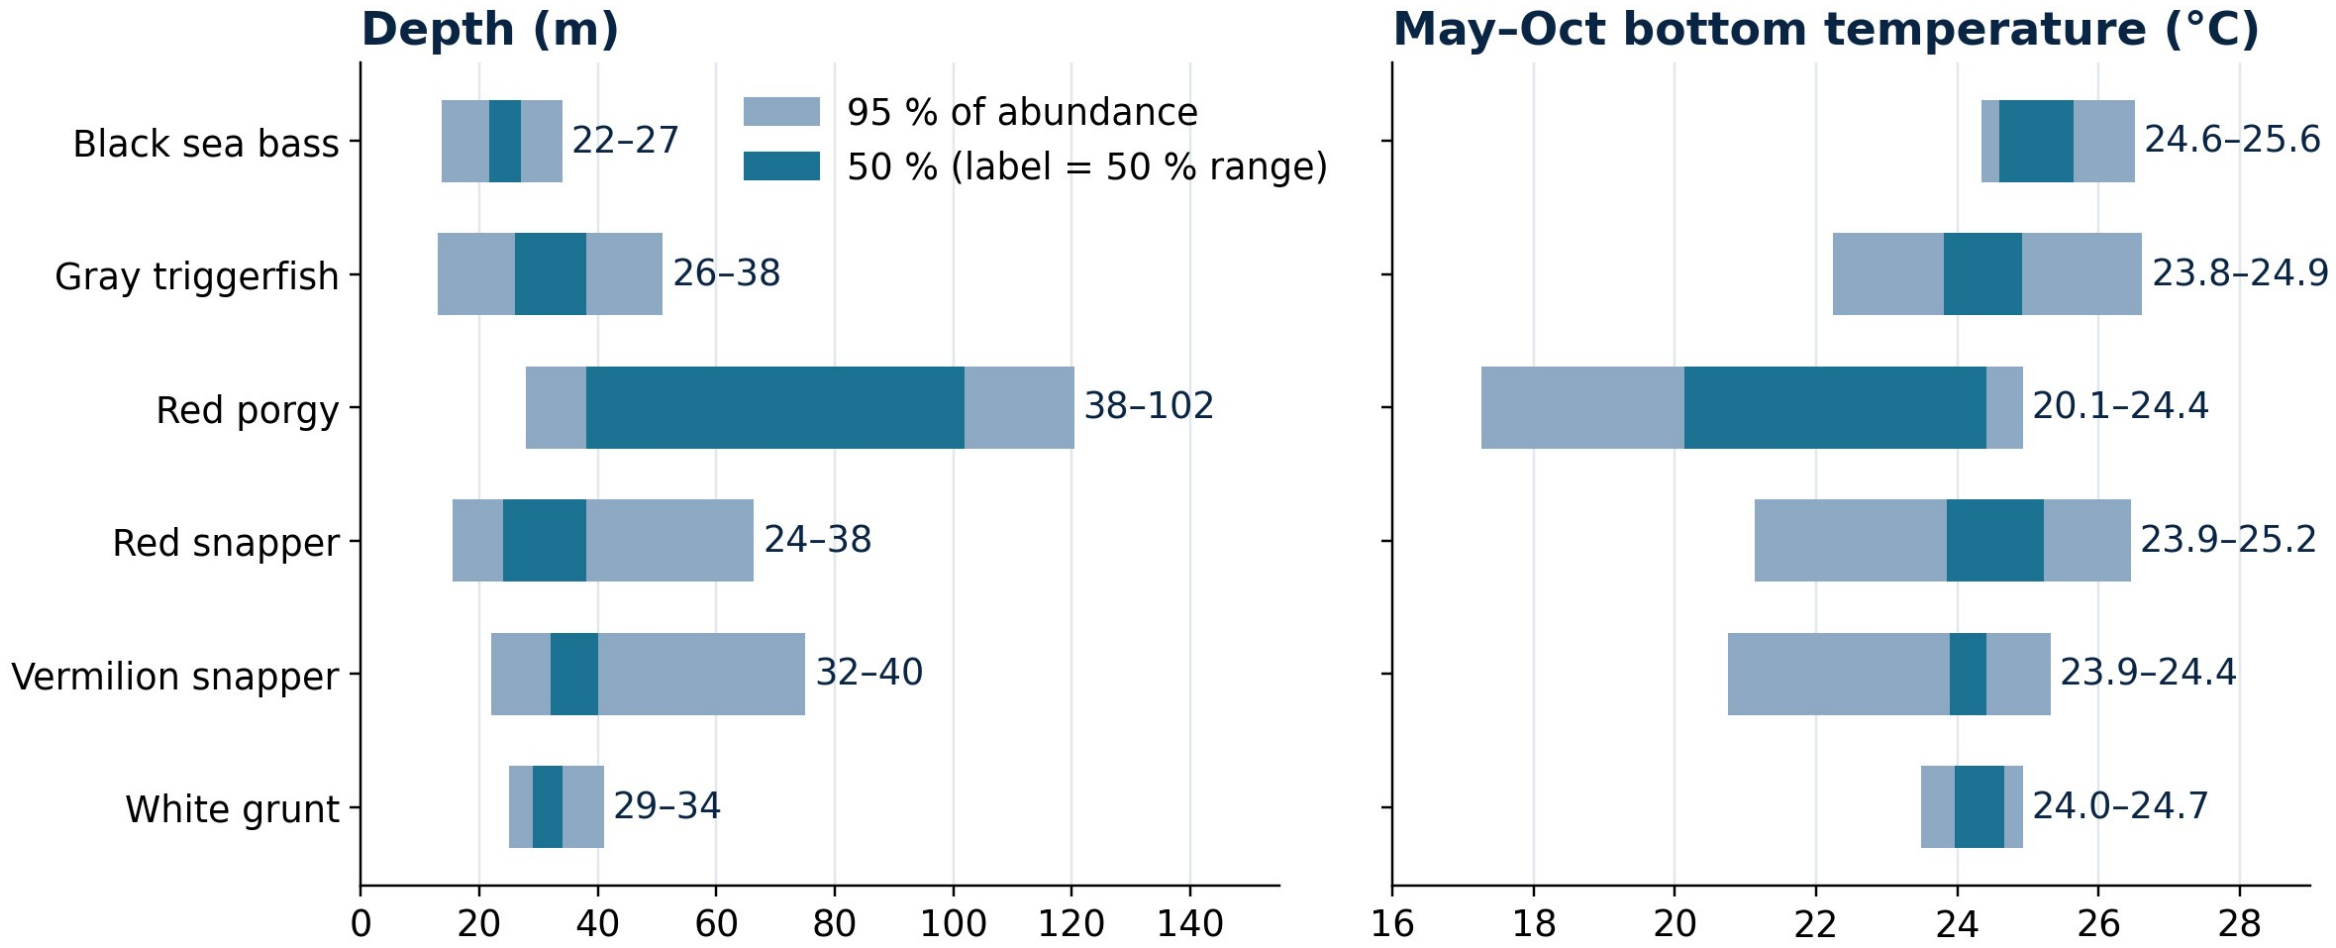
White grunt
mean inside tier = 0.7



NC STATE UNIVERSITY

Depth and temperature envelopes for the six species

Density-weighted ranges holding 95 % (light) and 50 % (dark) of persistent abundance



Depth from the survey grid; bottom temperature is the GLORYS12 May–October mean. Labels give the 50 % range.

NC STATE UNIVERSITY

Habitat area and environmental envelopes

Per species: tier areas and the depth and temperature ranges holding 95 % and 50 % of persistent abundance

Species	95 % area (km²)	50 % area (km²)	25 % area (km²)	Depth 95 % (m)	Depth 50 % (m)	Temp 95 % (°C)
Black sea bass	31,625	625	100	14–34	22–27	24.3–26.5
Gray triggerfish	72,600	7,575	500	13–51	26–38	22.2–26.6
Red porgy	14,000	475	50	28–120	38–102	17.3–24.9
Red snapper	45,175	4,375	750	15–66	24–38	21.1–26.5
Vermilion snapper	31,325	400	125	22–75	32–40	20.7–25.3
White grunt	4,775	75	25	25–41	29–34	23.5–24.9

Comparing with current EFH designations

Current designation (Snapper-Grouper FMP)

- One habitat-type description for the whole complex: coral reefs, live/hard bottom, SAV, artificial reefs and outcroppings
- From shore to at least 600 ft (2,000 ft for wreckfish), where water temperature sustains adults
- Species-specific depth bands only for tilefishes (HAPCs)

What the model adds

- A quantitative, species-level surface for each modelled species
- Depth and temperature envelopes and tier areas in the terms EFH is written in
- Shows how much between-species variation the single description absorbs, and which species it fits well or poorly

Part II

Multispecies model

Eighteen species · borrowing spatial structure from the community

NC STATE UNIVERSITY

Why a multispecies model

Species share the same reefs, ledges and shelf: their spatial patterns are not independent

Rich: black sea bass

13,000+

detections

Its own data can map it, track it and estimate trends with confidence. A self-sufficient dataset.

Poor: gray snapper

~1,500

detections

Alone, maps are noisy and trends are guesses. Not enough signal to estimate a reliable spatial field.

Very rare: sand tilefish

763

detections

Only two of them in traps, so nothing before 2010; spread over 300 km. Modelling this species alone produces noise, not a map.

The key insight. Species share habitat, so a few community-wide spatial patterns, estimated from all species together, can carry the map for the ones with too few detections to estimate their own.

The 18 species

Ordered by positive samples (trap + video, 1990–2024); the six with their own single-species models are marked

Black sea bass *Centropristis striata* **13,169**	**Almaco jack** *Seriola rivoliana* 4,991	**Lane snapper** *Lutjanus synagris* 942
Gray triggerfish *Balistes capriscus* **12,448**	**Greater amberjack** *Seriola dumerili* 3,320	**Red grouper** *Epinephelus morio* 927
Red porgy *Pagrus pagrus* **12,190**	**Scamp** *Mycteroperca phenax* 3,175	**Mutton snapper** *Lutjanus analis* 817
Vermilion snapper *Rhomboplites aurorubens* **11,113**	**Gag** *Mycteroperca microlepis* 1,593	**Sand tilefish** *Malacanthus plumieri* 763
Red snapper *Lutjanus campechanus* **8,045**	**Gray snapper** *Lutjanus griseus* 1,470	**Banded rudderfish** *Seriola zonata* 711
White grunt *Haemulon plumierii* **5,783**	**Hogfish** *Lachnolaimus maximus* 986	**Graysby** *Cephalopholis cruentata* 621

own single-species model (Part I) multispecies model only

Previous work: a multispecies, video-only model

Cao, Craig & Damiano (2024, Ecosphere): spatiotemporal dynamics of Atlantic reef fishes off the southeastern U.S. coast

Received: 23 April 2023 | Revised: 31 January 2024 | Accepted: 27 February 2024

DOI: 10.1002/ecs2.4868

ECOSPHERE
AN ESA OPEN ACCESS JOURNAL

ARTICLE

Coastal and Marine Ecology

Spatiotemporal dynamics of Atlantic reef fishes off the southeastern U.S. coast

Jie Cao[1] | J. Kevin Craig[2] | Matthew D. Damiano[1]

The baseline model

$$\log \mu_{s,g}(x,t) = \alpha_{s,g} + \gamma_t + \log E(x,t) + \omega_{s,g}(x) + \varepsilon_{s,g}(x,t)$$

$$Y_{s,g}(x,t) \sim \text{NegBin}(\ \mu_{s,g}(x,t),\ \varphi_{s,g}\ )$$

$\alpha_{s,g}$: mean catch level for species s and gear g

γ_t : year effect in year t

$E(x,t)$: sampling effort (log-offset)

$\omega_{s,g}(x)$: persistent spatial field for species-gear pair (s,g)

$\varepsilon_{s,g}(x,t)$: year-specific spatial anomaly field

What is missing

- Nothing links the two gears except the year effect
- Nothing links species: a rare species must estimate its own ω and ε

Two lenses on one map, plus a trap-only field

The same two-gear logic as Part I, written for every species at once

The shared-field model

$$\omega_{s,\text{video}}(x) = W_s(x)$$

$$\omega_{s,\text{trap}}(x) = \gamma_s \, W_s(x)$$

- Gears do not see different populations: each species has one population map W_s(x)
- Video is the reference lens; the trap sees the same map through a scaling factor γ_s

The extended trap model

$$\omega_{s,\text{trap}}(x) = \gamma_s \, W_s(x) + \lambda_s \, Q(x)$$

- Q(x): one trap-only spatial field, where traps fish well, shared by all species
- λ_s: each species' sensitivity to it
- Captures habitat selectivity the gears have in common, separately from the species' distributions

Species share the habitat

Community factors are estimated from all species; a rare species estimates three loadings, not a whole map

$$W_s(x) = \Lambda_{s1}\, F_1(x) + \Lambda_{s2}\, F_2(x) + \Lambda_{s3}\, F_3(x) + \delta_s(x)$$

F — Community spatial patterns

A few shared fields estimated from all species jointly; the same construction in space-time with G_m(x,t)

Λ — Species loadings

How much of each community pattern a species carries: an ordination of the community

δ_s — What is uniquely its own

A species-specific residual field with its own variance σ_s

Automatic, data-driven regularisation. If a species has private spatial signal, δ_s carries it; if not, σ_s shrinks toward zero and the species falls back on the community blend.

The assembled model

Community factors → species fields → two gear lenses → negative-binomial counts

1 Community

$$F_1(x),\ \ldots,\ F_L(x) \qquad G_1(x,t),\ \ldots,\ G_M(x,t)$$

A few shared spatial patterns and year-anomaly fields, estimated jointly from all species

2 Species

$$W_s(x) = \Sigma_\ell\ \Lambda_{s\ell}\ F_\ell(x) + \delta_s(x)$$
$$E_s(x,t) = \Sigma_m\ \Psi_{sm}\ G_m(x,t) + \eta_s(x,t)$$

Blend of community factors + private residual, in space and in space-time

3 Gears

Video: $W_s(x) + E_s(x,t) + T_s(t)$

Trap: $\gamma_s\ W_s(x) + \beta_s\ E_s(x,t) + \lambda_s\ Q(x) + t_s\ T_s(t)$

Video is the reference; the trap adds γ_s, β_s, λ_s Q and its own year scaling t_s

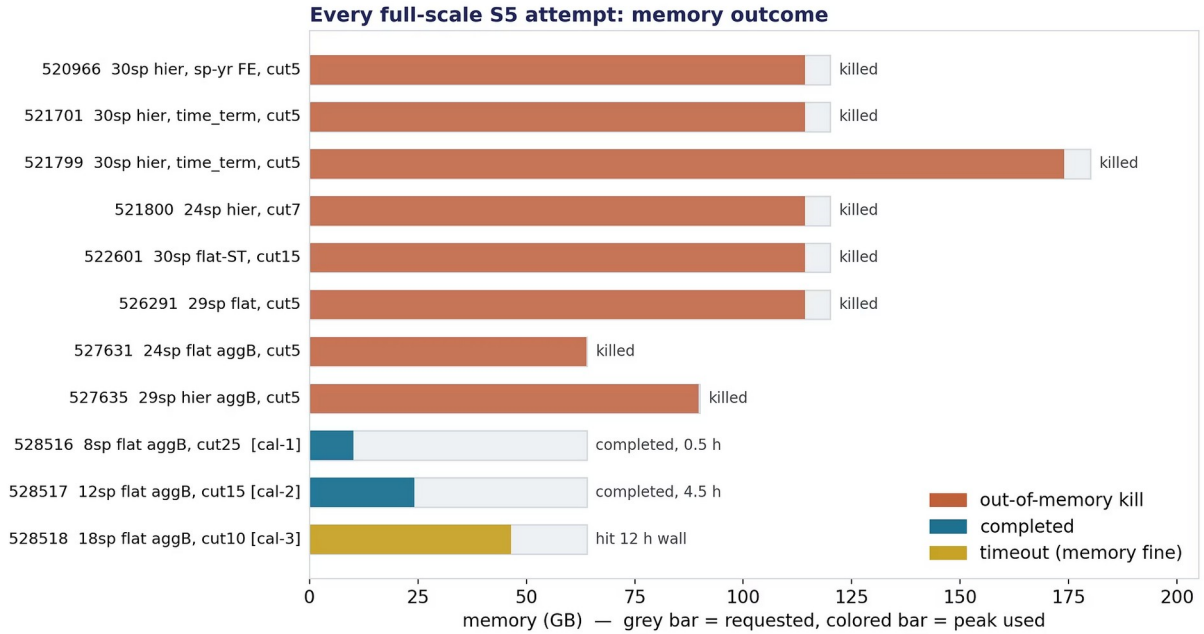
4 Observations

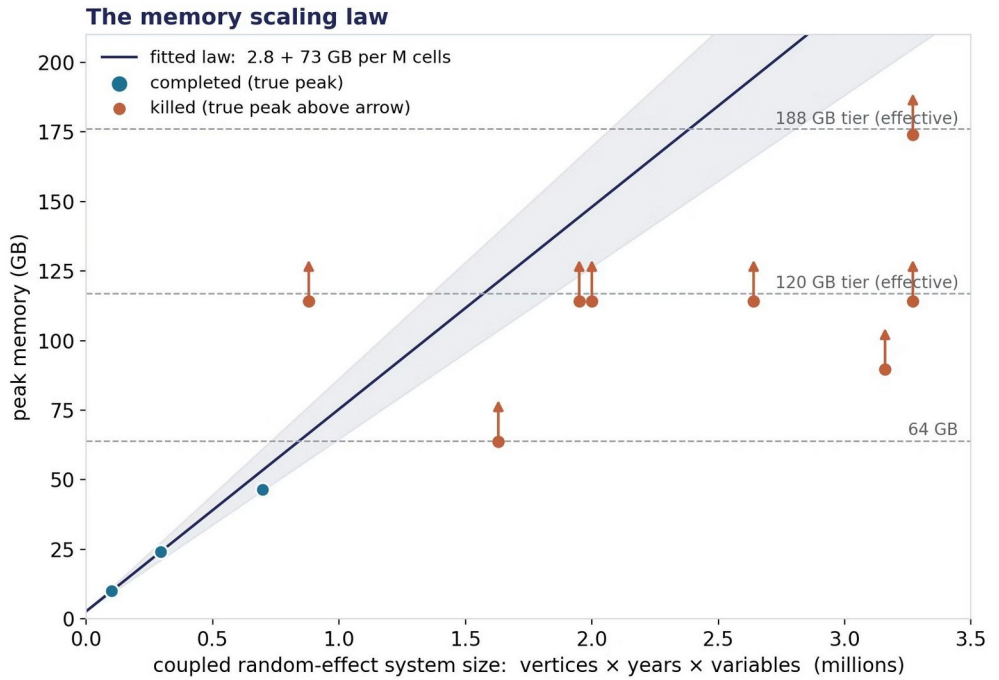
$$Y_{s,g}(x,t) \sim \mathrm{NB2}(\ \exp\{\alpha_{s,g} + \ldots\} \cdot \mathrm{effort}(x,t),\ \varphi_{s,g}\ )$$

Negative binomial with a species × gear dispersion parameter; effort as an offset

Fitting at scale: memory, not time, was the constraint

Every full-scale attempt was killed at the cluster memory cap; peak memory scales with vertices × years × variables





The scaling law

peak ≈ 3 GB + 73 GB per million random-effect cells

The fix

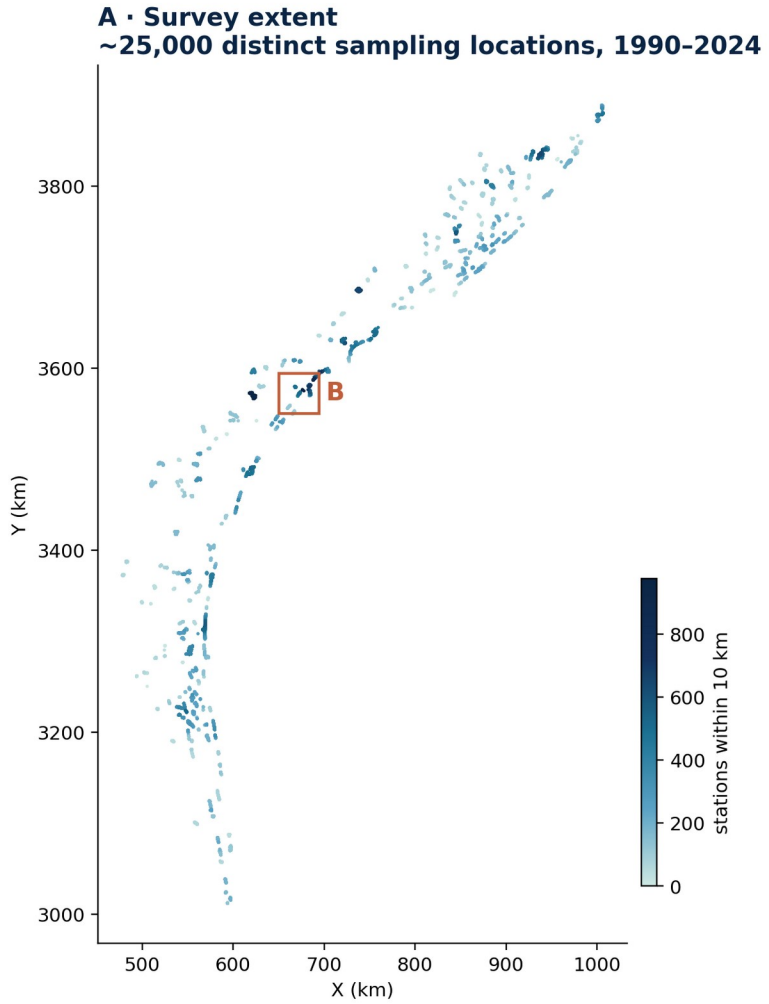
Aggregate traps within a sampling event (5-km cutoff): 1.30 M → 0.2 M rows; no spatial signal below ~20 km

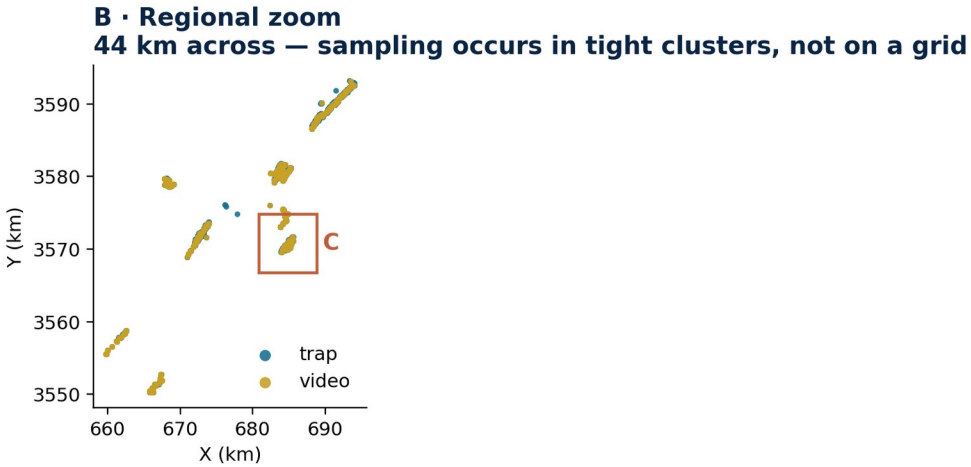
Production fit

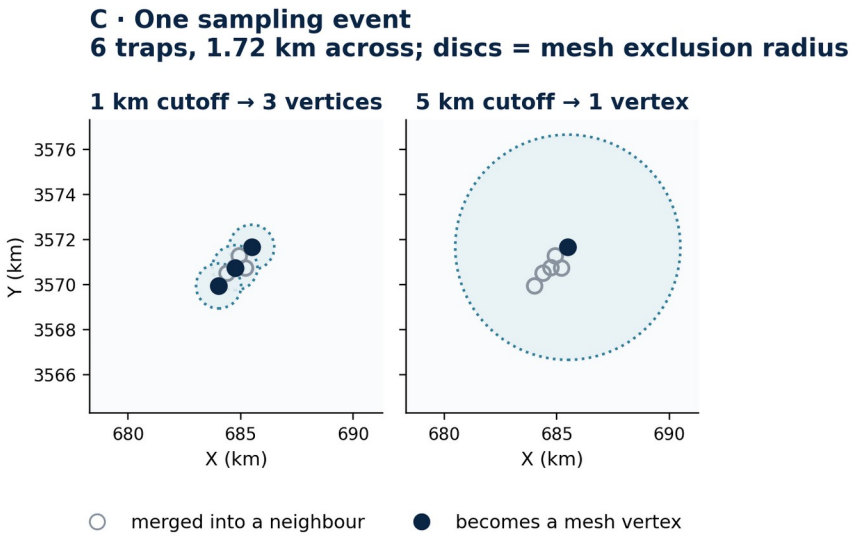
18 species, 15-km mesh, ~0.8 M cells: 47 GB, ~15 h

NC STATE UNIVERSITY

Spatial variation vs observation error?







NC STATE UNIVERSITY

Structure selection: what was compared

18 species, identical data and mesh across all candidates

1 Flat vs hierarchical

Flat: a species map is a blend of the community factors only (Λ·F, no residual). Hierarchical: community factors plus a species-specific residual field δ_s with its own variance.

2 Free σ vs shared σ

Does each species get its own residual variance σ_s (34 extra parameters), or do all species share one?

3 Rank

How many community factors: spatial rank 2 to 8, spatiotemporal rank 1 to 3. Candidates are judged on AIC and on whether every loading is identified.

4 Gear structure

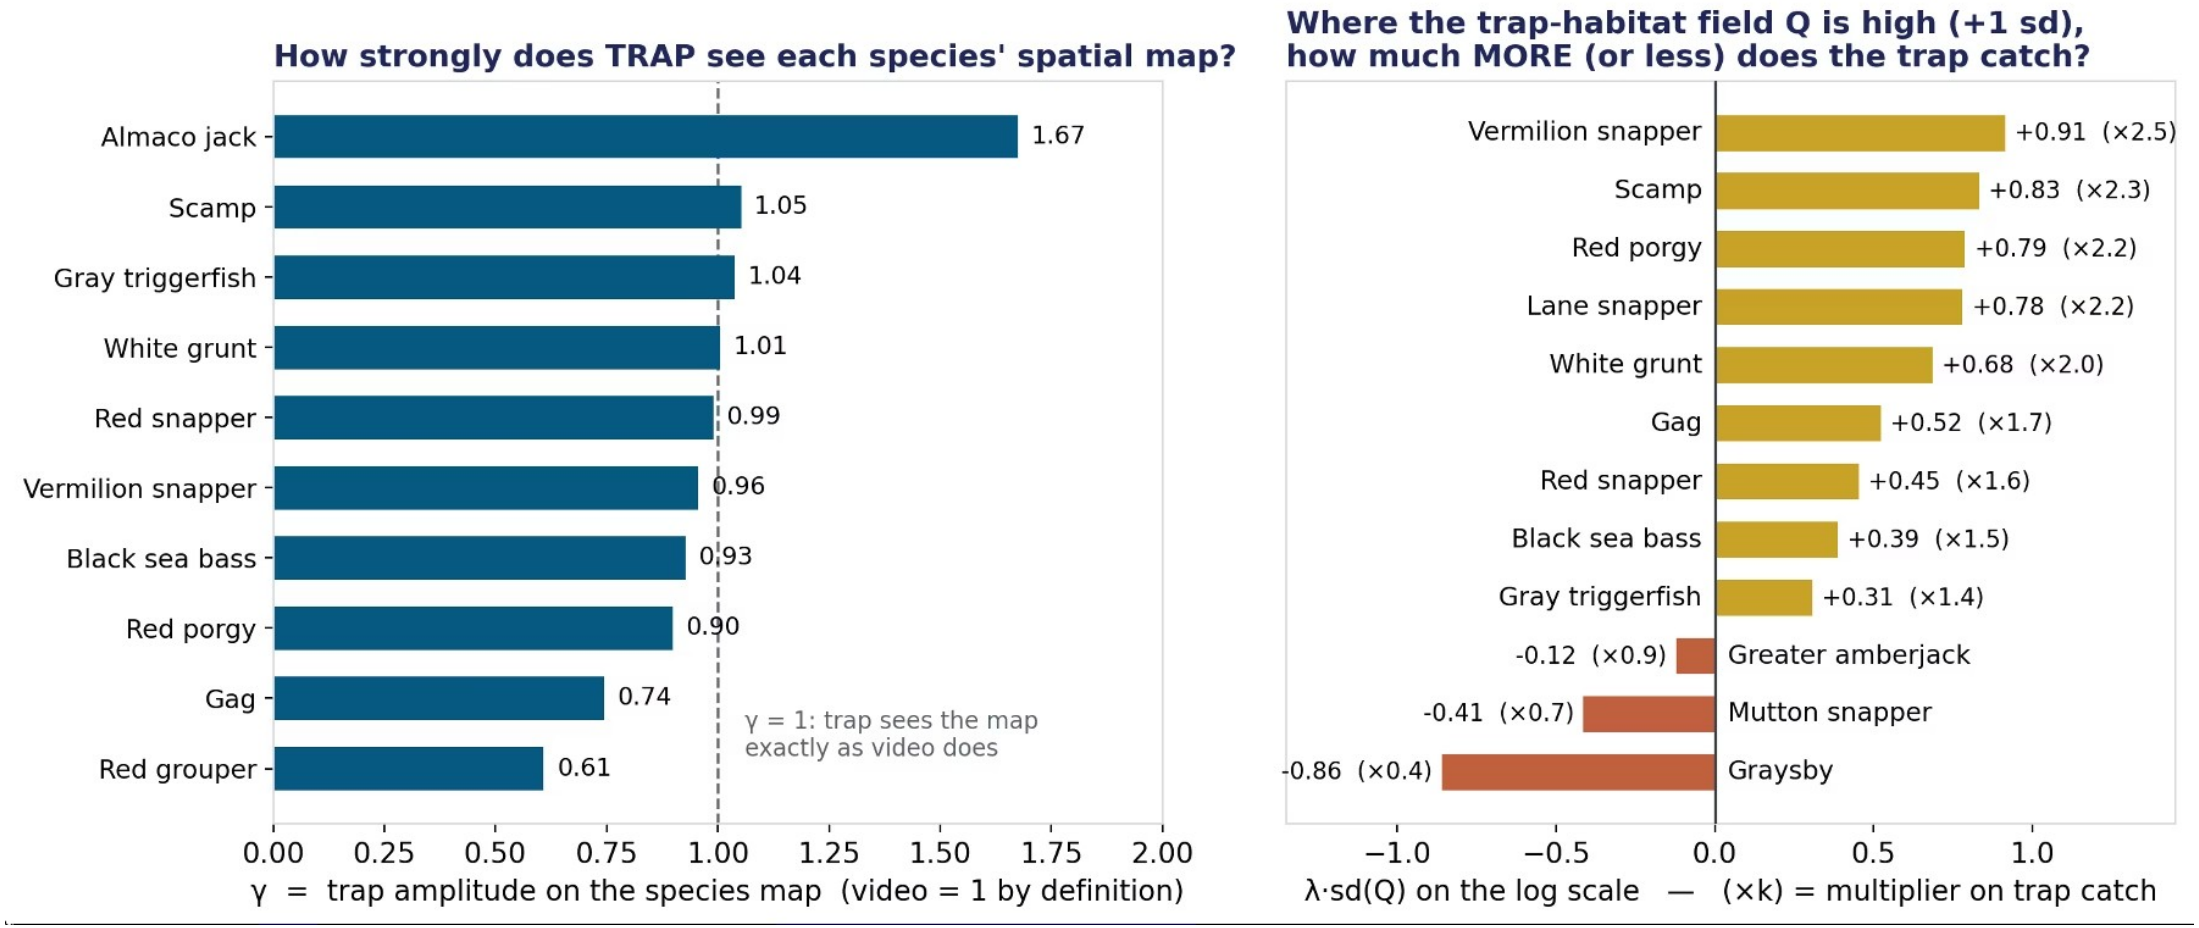
The two-gear treatment from Part I, plus the shared trap-only field Q with a species-specific sensitivity λ_s.

Structure selection: AIC versus identifiability

Six alternatives on identical data and mesh, 18 species

Structure	AIC	ΔAIC	npar	Status
Hierarchical, species σ, rank 5/3	284,124	0	352	lowest AIC — not converged
Hierarchical, species σ, rank 3/2	**284,356**	**232**	**307**	selected: highest converging rank
Hierarchical, species σ, rank 2/1	284,575	451	310	
Hierarchical, shared σ, rank 3/2	284,623	499	309	
Flat, shared σ, rank 8/3	284,663	539	550	
Flat, shared σ, rank 4/2	285,164	1,040	394	

Gear findings

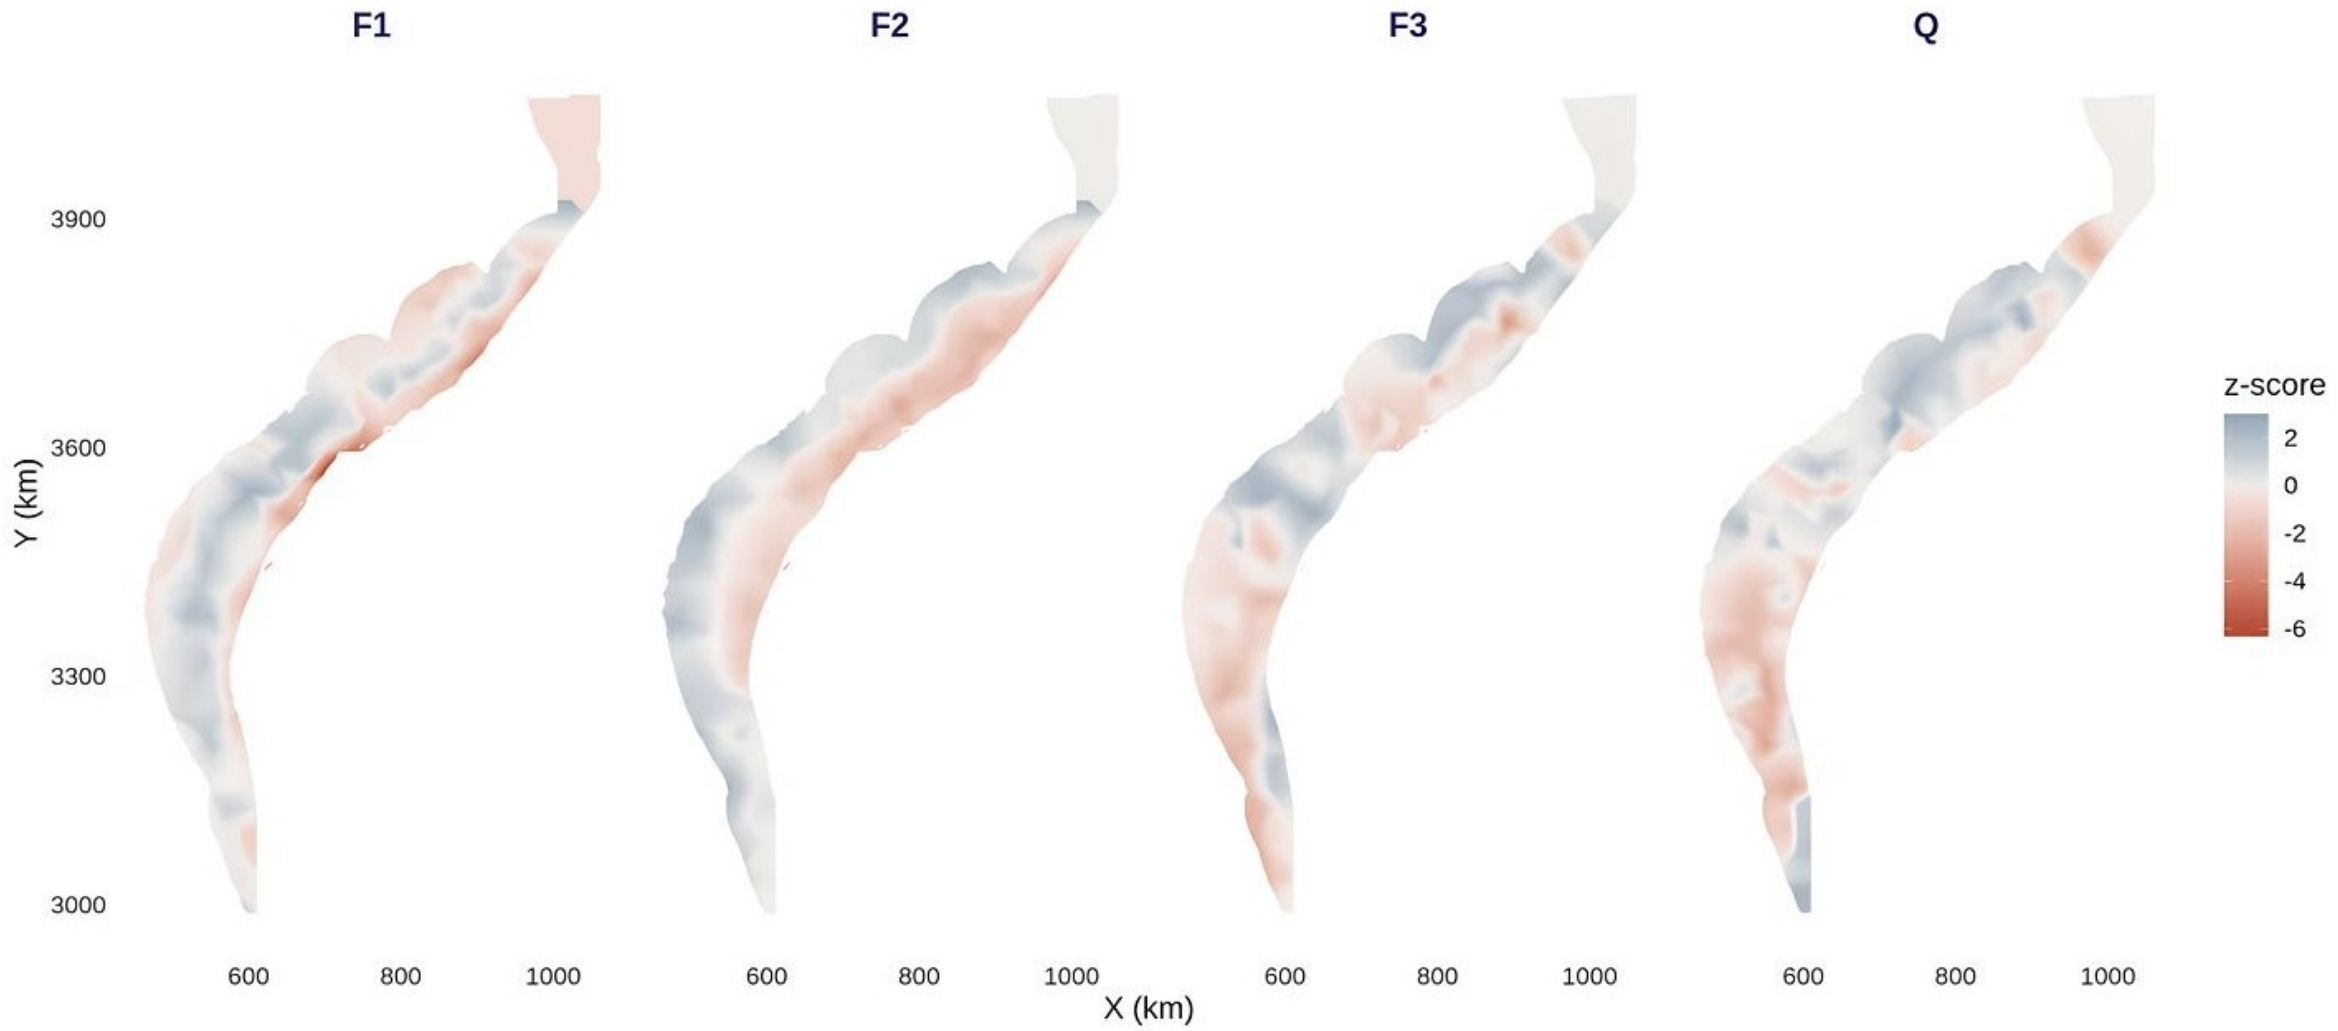


Reading. γ near 1 for most species: the trap sees the same map as the video. Q raises trap catch 1.4–2.5× for most species and lowers it for graysby, mutton snapper and greater amberjack.

NC STATE UNIVERSITY

Community spatial patterns and the catchability field

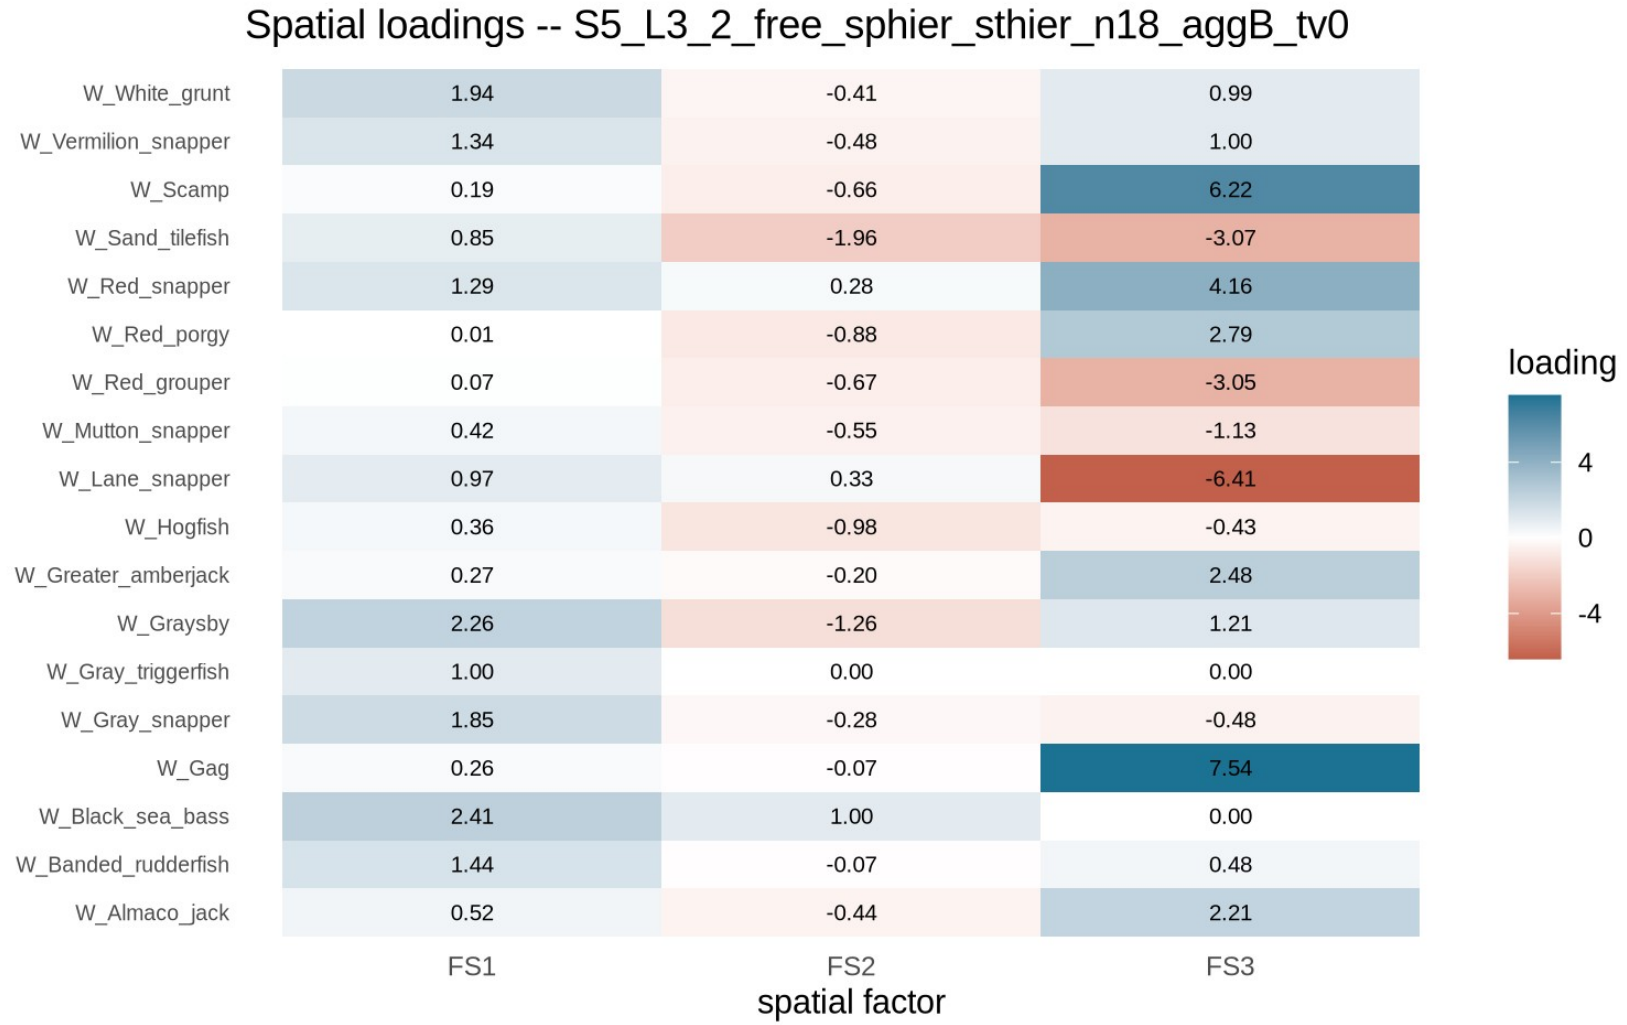
F1–F3 are the habitat axes every species blends; Q is the trap-only field, shared by all species



NC STATE UNIVERSITY

Community ordination: who loads on which patterns

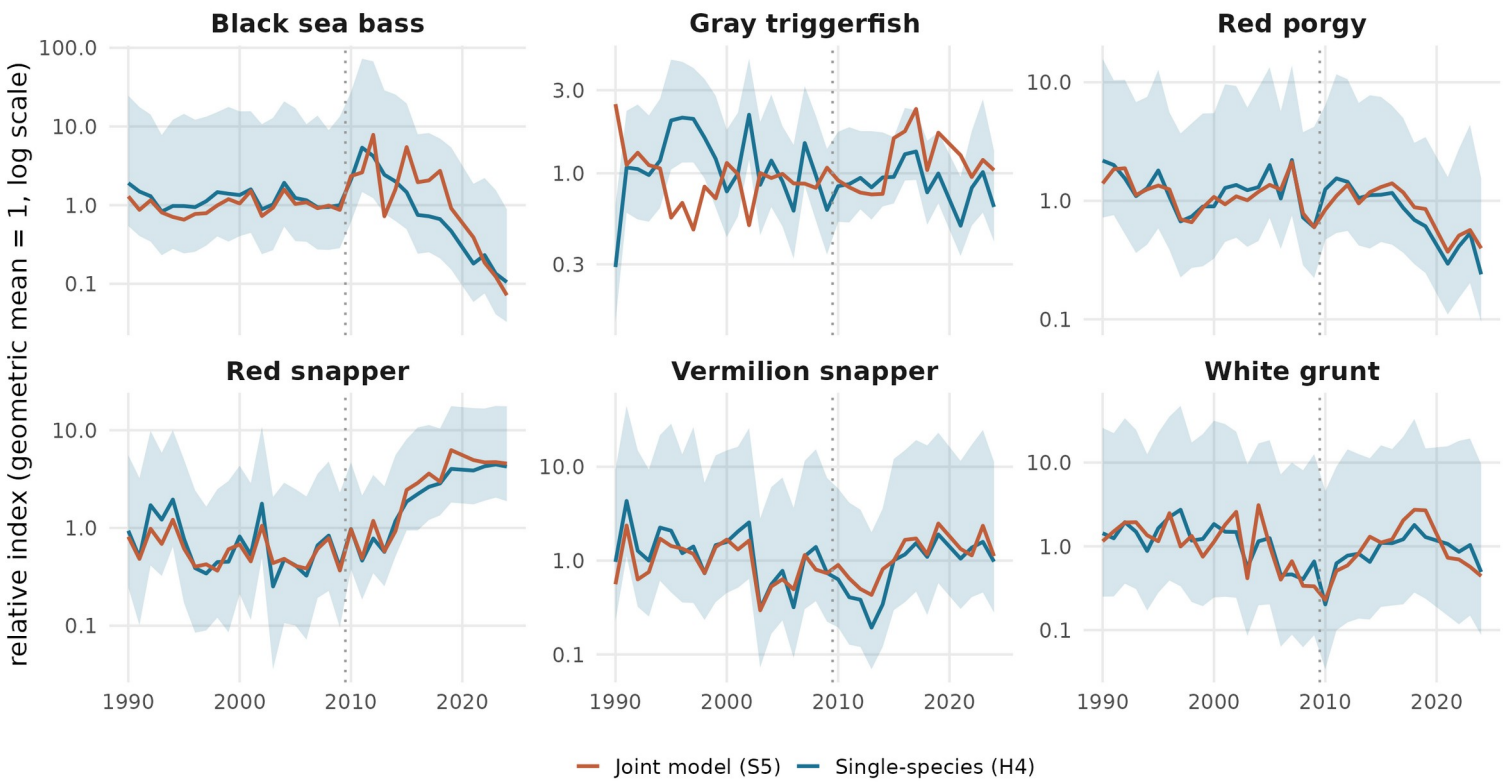
Species loadings on the three spatial factors, rank 3/2 fit



- FS1 positive for nearly all species: the shared reef axis
- FS3 is the main contrast: gag, scamp, red snapper, red porgy (+) vs lane snapper, red grouper, sand tilefish (–)
- Similar rows = shared spatial niche

Joint vs single-species models: do they agree?

Annual index for the six species with both models; band = single-species 95 % interval



Species	Map r	Index r	Index r 2010+
Black sea bass	0.77	0.79	0.83
Gray triggerfish	0.63	**−0.54**	0.44
Red porgy	0.79	0.91	0.92
Red snapper	0.81	0.95	0.98
Vermilion snapper	0.74	0.84	0.92
White grunt	0.87	0.78	0.83

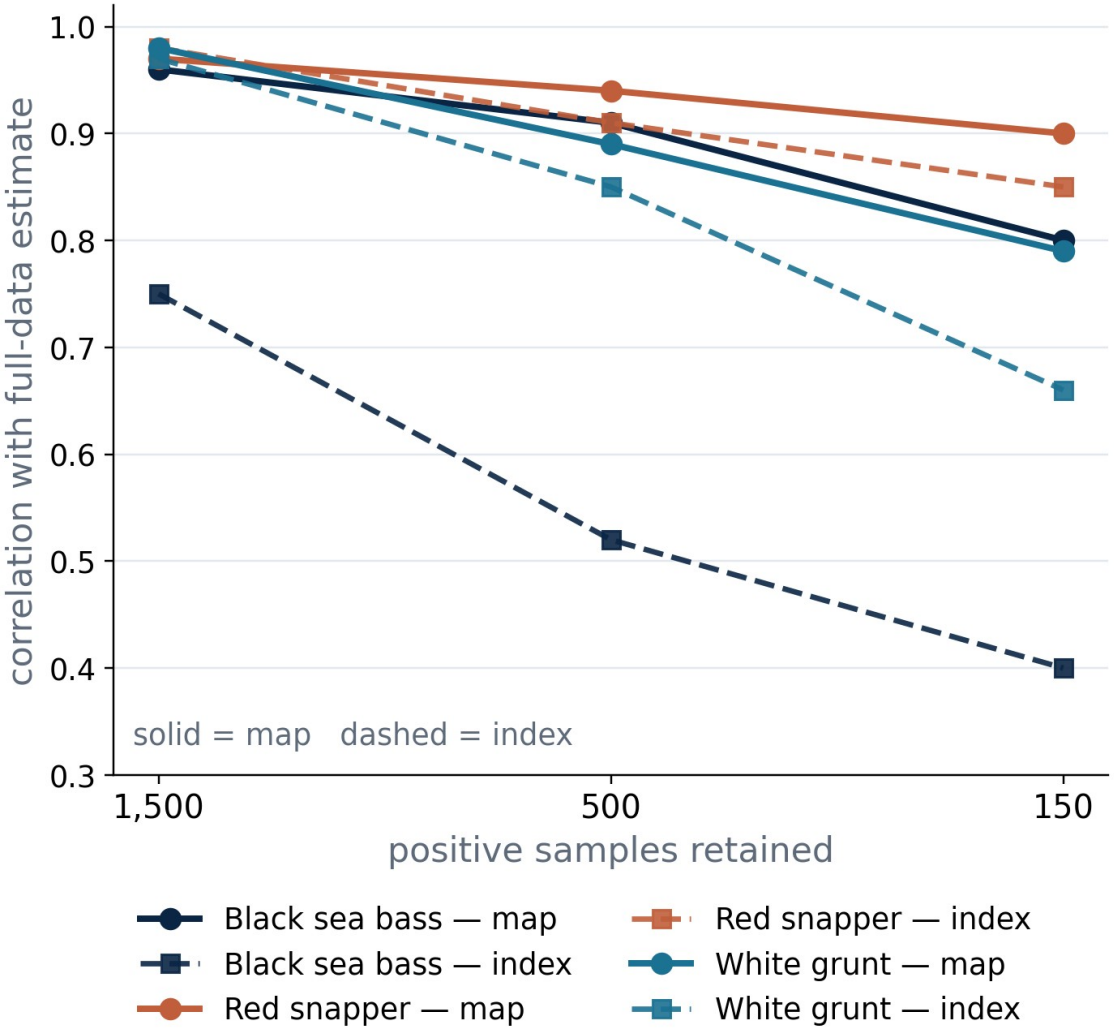
Map r: year-averaged log density on matched cells (joint model 15-km mesh vs 5-km, so a floor). Index r: centred log index; joint estimates fall inside the single-species interval in 77–100 % of years.

- Maps agree (r 0.63–0.87, 4,000 cells)
- Indices agree for five species; gray triggerfish disagrees before 2010
- Joint-model COG swings 2–3× more than single-species

NC STATE UNIVERSITY

Validation: recovering a map from few detections

Three well-sampled species thinned to 1,500, 500 and 150 positives and refitted



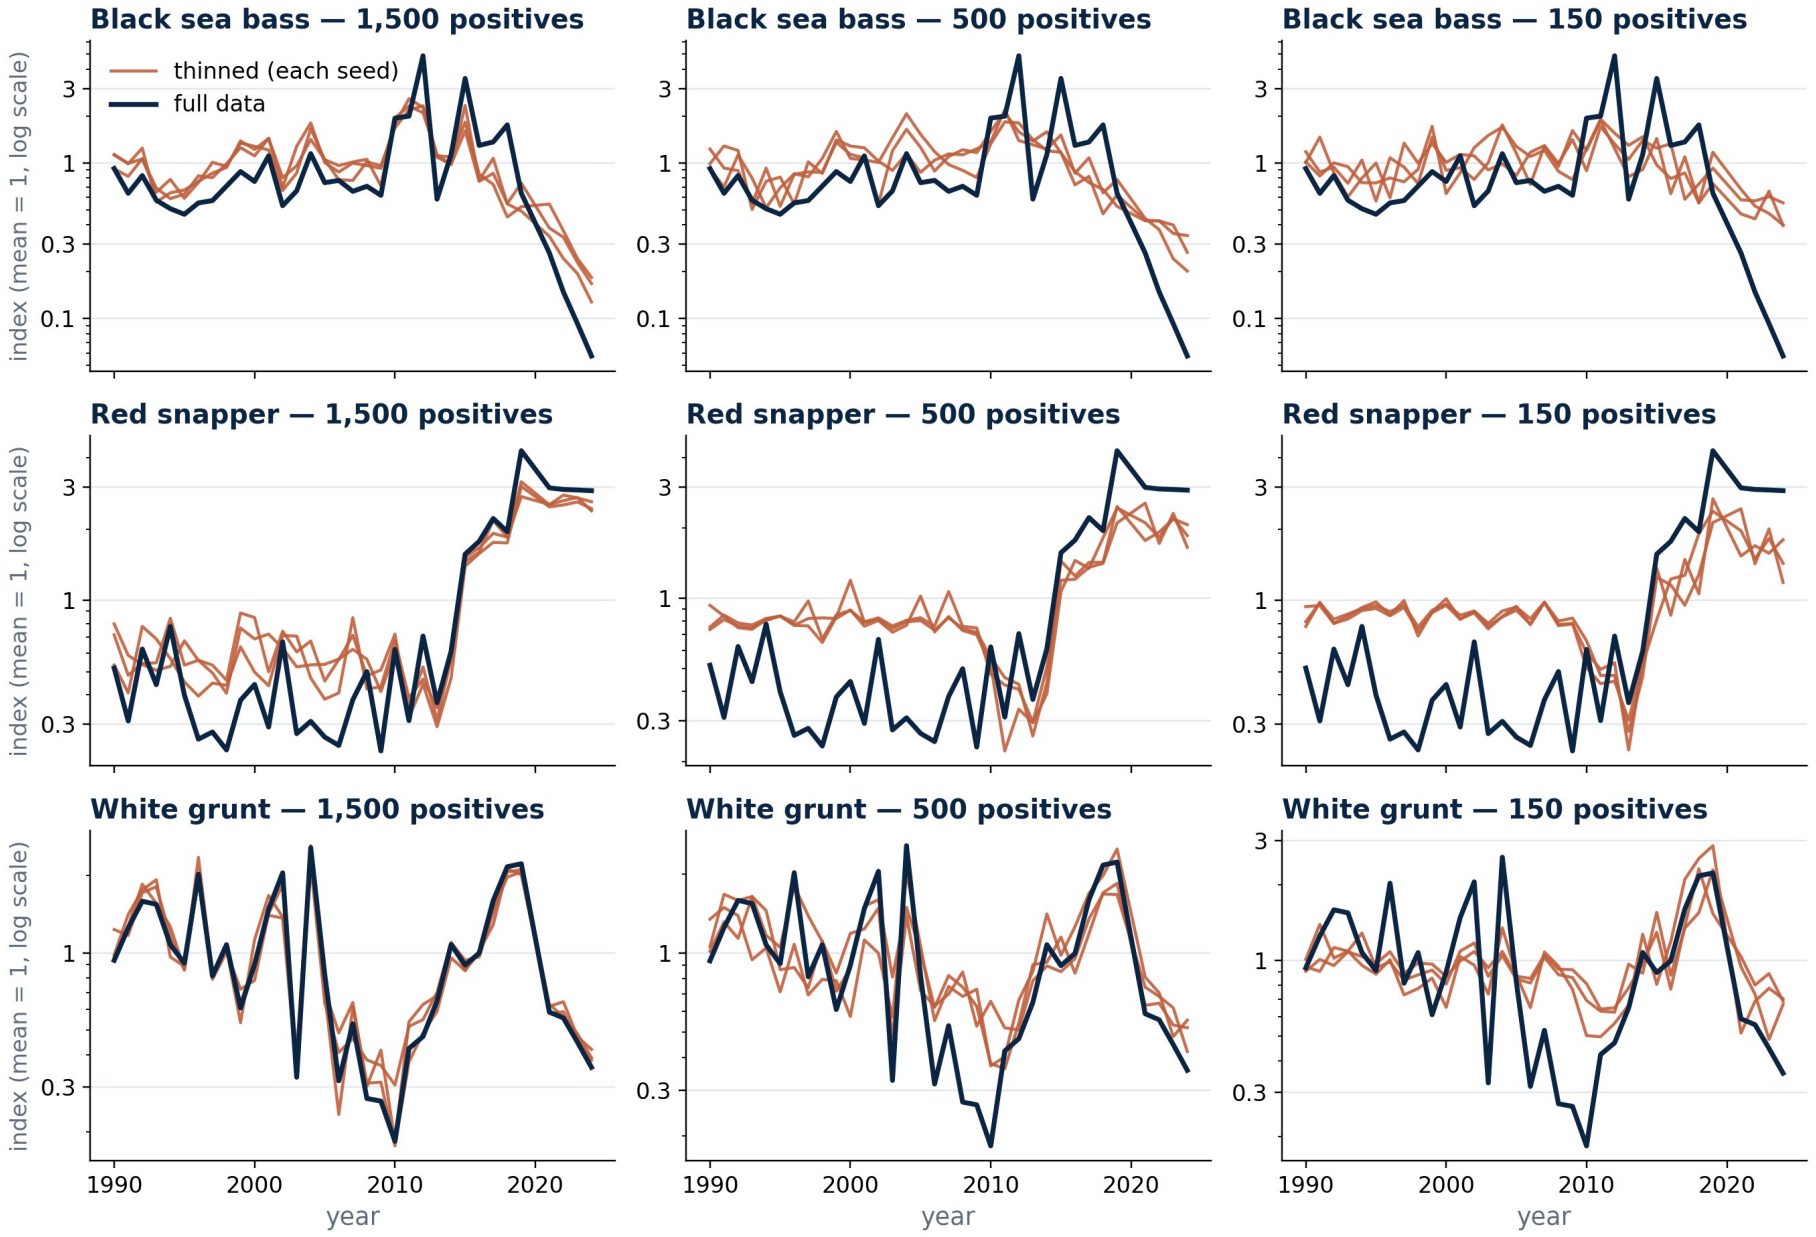
Maps recover

Even at 150 positives, the rarest tier, the thinned map still matches the full-data map at r = 0.79–0.90: the community factors supply the pattern; the sparse data only set the loadings

Indices recover when a species tracks the community

Red snapper 0.98 → 0.85, white grunt 0.97 → 0.66 as positives fall from 1,500 to 150: the community annual factors carry most of their trend

NC STATE UNIVERSITY

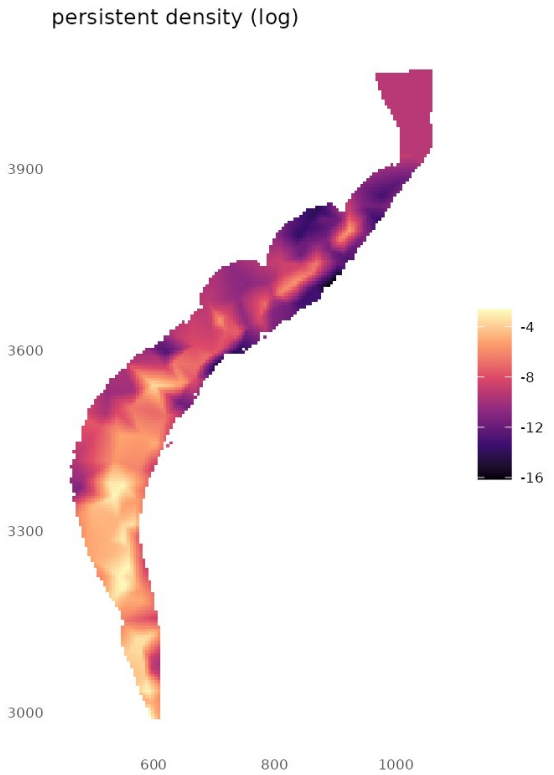


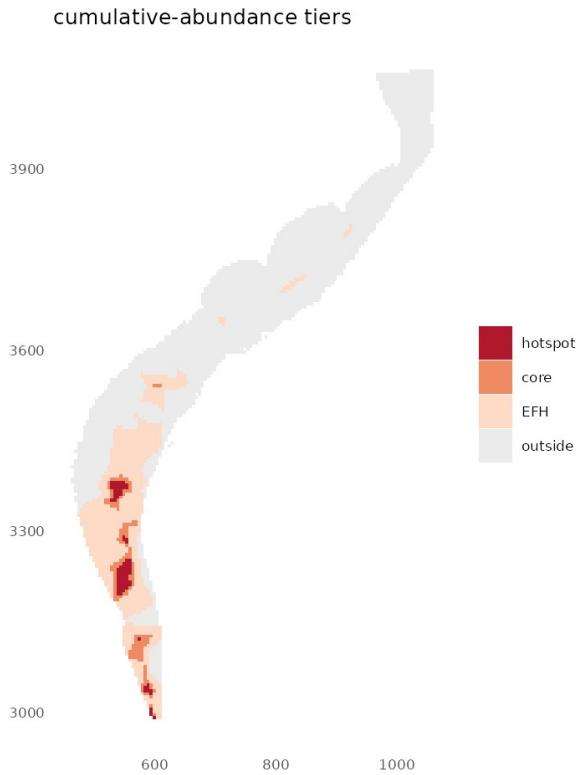
- Annual index on a log scale: full data (navy) vs thinned to 1,500, 500 and 150 positives (red, three seeds each)
- Full-data positives (trap + video): black sea bass 13,169; red snapper 8,045; white grunt 5,783

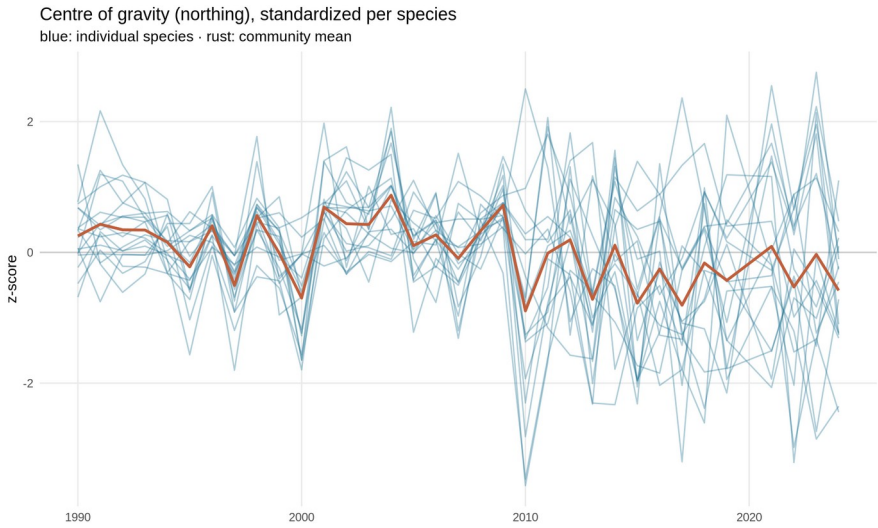
NC STATE UNIVERSITY

Multispecies products (preliminary) for 18 species

Gray snapper: 46 trap and 1,424 video positives — not modellable alone







Persistent log density · EFH tiers (95 / 50 / 25 %) · centre of gravity for all 18 species with the community mean.

NC STATE UNIVERSITY

Summary

What we can say now

- Trap and video can be combined: the shared-field structure (H4) is identifiable, is supported by both real data and simulation, and predicts across gears
- Six species have production models: indices 1990–2024, distribution indicators and EFH products, robust to era, gear and covariate choices
- Depth and temperature are the habitat drivers; hard bottom adds nothing at this resolution; every species now occupies warmer water than in 1990, with no northward or depth shift
- The multispecies model recovers maps for data-limited species from the community

Caveats we carry forward

- Vermilion snapper covariates weakly supported; red snapper uses the null covariate model
- Joint model: gray triggerfish trap year-loading estimated negative; disagrees with its single-species index before 2010
- Video-only species carry no information before 2010
- Community factors and loadings are shown unrotated; a PCA rotation would make the factors easier to interpret without changing the fit
- Multispecies products are preliminary

Acknowledgements

Funding
South Atlantic Fishery Management Council, contract SA-25-111, under NOAA's Inflation Reduction Act climate award

Data
Southeast Reef Fish Survey (SERFS): SCDNR MARMAP, SEAMAP-SA and NOAA SEFIS trap and video programs; GLORYS12 reanalysis (Copernicus Marine Service)

Collaborators
Kevin Craig (NOAA SEFSC); Janet Nye and Jade Hong (UNC Institute of Marine Sciences); the project oversight team and SAFMC staff

Computing and software
NC State University High Performance Computing (Hazel); tinyVAST (Thorson et al. 2024), R-INLA mesh tools

Thank you

jcao22@ncsu.edu

NC STATE UNIVERSITY