What Race and Ethnicity Hide

Bayesian Approximate Computation with Hierarchical Framework (BACH) for Subgroup Voting Behavior under Partisan Polarization, Cross-Cutting Cleavages, and Survey Nonresponse

Tyler Rongxuan Chen | Ph.D. Candidate, Political Science and Scientific Computing
tylercr@umich.edu

Motivation: Three Obstacles

Pre-election surveys in 2024 disagreed by **more than 20 points** on the Hispanic vote, and AAPOR reports polls have overstated Democratic support for **three straight cycles** — a bias no weighting recipe removed. The residual error lives **within demographic cells**.

1 · Survey nonresponse. Response propensity correlates with vote choice within cells (Tyler et al. 2026): the within-cell representativeness that poststratification relies on fails.

2 · Partisan polarization. Cell-level vote shares concentrate near 0 and 1, widening the identified set exactly where parametric EI methods strain.

3 · Cross-cutting cleavages. Cuban- and Dominican-American voters differed by **20+ points** in 2020; the White college gap approaches the Black–White gap in many districts. Race-level estimators absorb this variation into a residual.

The BACH Framework

Rather than enriching covariates around a race-level estimator, BACH **cuts the population cell structure itself**: inference runs over within-race cells from the outset, disciplined by aggregate returns.

The ecological identity is underdetermined — C unknowns, one constraint per contest:

$$Y_{ie} = \sum_{c=1}^{C} s_{ic}\,\theta_{ic}^{(e)} + \varepsilon_{ie}$$

Survey-calibrated Beta priors enter as **soft centers, not point estimates** (N_0 set far below survey n):

$$\theta_{ic}^{(e)} \sim \mathrm{Beta}\big(\bar{\theta}_c^{(e)} N_0,\ (1-\bar{\theta}_c^{(e)}) N_0\big)$$

Weight K ≈ 100,000 forward-simulated particles by a **continuous ABC kernel** (SMC in the full pipeline):

$$w_i^{(k)} \propto \exp\left(-\sum_{e=1}^{E_i} \frac{(\hat{Y}_{ie}^{(k)} - Y_{ie})^2}{2h^2}\right)$$

Obstacle	BACH's answer
Nonresponse	Loose survey prior — a soft center the aggregates update
Polarization	Simulation-based weighting — no parametric family at the 0/1 boundary
Cleavages	Within-race cells are the objects of inference from the start

Multi-contest and differenced designs tighten the identified set; IPF recovers joint tables where the census publishes only margins — so the pipeline **travels to comparative settings** with limited individual-level data.

In progress — Study 3: NMAR. A third validation study, targeting nonignorable (NMAR) nonresponse directly, is under construction — ask me about the design.

Validation I: Simulation

375 conditions (3 DGPs × 5 segregation × 5 polarization × 5 nonresponse levels), **200 replications each** — nonresponse up to near-complete suppression of one party within cells. 128 cells, ~7M voters per replication.

Method	T1 MAE	T1 Cov.	T2 MAE	T2 Cov.	Time (s)
BACH	0.037	0.994	0.026	0.963	241
MRP	0.034	0.931	0.025	0.889	380
GQ	0.066	0.691	—	—	4,175
GQ Expanded	0.047	0.968	—	—	2,890
SEINE	0.080	0.764	0.255	0.624	0.1
King	0.110	0.903	—	—	0.2

Tier 1 = race-level shares; Tier 2 = within-race gaps; 95% nominal coverage.

BACH is the **only method tested with nominal coverage at every tier**, at point accuracy indistinguishable from MRP — and it is **direction-agnostic**: reversing which party is suppressed raises MRP's Race-2 error by 78%; BACH's coverage exceeds 99% either way.

A substantive analysis needs **all nine quantities right at once** — and there MRP drops to a coin flip: joint coverage **41.3% vs. BACH's 82.4%**.

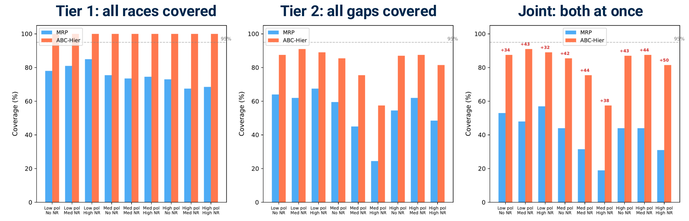

Joint Tier-1 × Tier-2 coverage (ABC-Hier = BACH). At high polarization × high nonresponse, MRP's largest-group coverage falls to 79%; BACH stays above 97%.

Validation II: Florida Ground Truth

The L2 Florida voter file is a **census of registrants**: cell-level Democratic registration is observed exactly. I aggregate to 67 counties, discard the microdata, and recover **32 race × gender × age cells** from aggregates alone.

Method	RMSE	MAE	Cov.	Out of [0,1]
BACH	0.019	0.015	0.969	0
MRP	0.167	0.114	0.656	0
MRP (calibrated)	0.147	0.101	0.656	0
SEINE	0.527	0.483	0.312	31

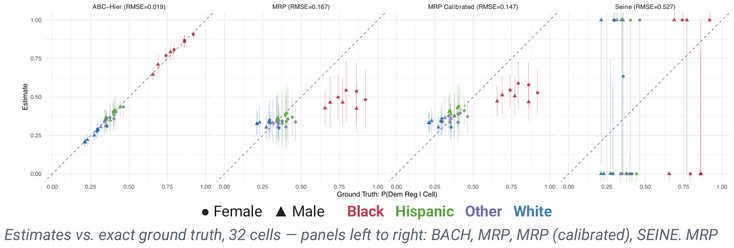

Estimates vs. exact ground truth, 32 cells — panels left to right: BACH, MRP, MRP (calibrated), SEINE. MRP compresses cells toward the center even after calibration; BACH errs by under 2 points per cell.

What Race-Level Estimates Hide

Kuriwaki et al. (APSR 2024) find racial polarization in 50 competitive districts predominantly geographic. An **asymmetric 9-cell BACH** — White and Black voters split by education, Hispanic voters by national origin — shows that story is incomplete:

≈ 25% of cell-level variance comes from within-race contrasts that a race-only decomposition assigns to the residual.

White voters: a 49-point education gap, positive in 96% of districts.

Hispanic voters: the near-zero aggregate education gap is a **Simpson's paradox** — national origin explains **7.3× more variation** than education, and the more-Democratic origin groups have lower college attainment.

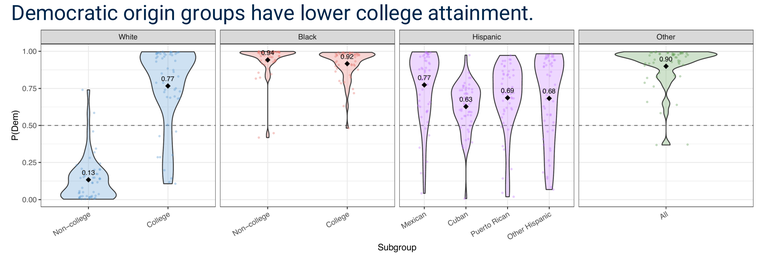

Within-race cleavages in P(Dem) pooled across 50 CDs: points are district posterior means, diamonds the grand mean.

Geographic variance, recovered: MRP reproduces 13% of actual within-CD precinct-level variance; contextual BACH reaches **80%** — from the same data.

Four More Replications

Study	Within-race decomposition finding
de Benedictis-Kessner & Warshaw (APSR 2020) · accountability	No wage-growth gradient in any race × education cell; White non-college −6.7 pp across all deciles.
Baccini & Weymouth (APSR 2021) · deindustrialization	White non-college −5.5 to −8.8 pp: concave, front-loaded dose–response.
Ferrer, Geyn & Thompson (APSR 2024) · election admin.	RD effect compresses the racial gap: White +2.1 pp vs. non-White +0.7 pp.
Morse et al. (AJPS 2025) · ballot design	Broward's excess undervote is nearly race-blind in absolute terms (+3.6 vs. +3.5 pp) — but it compounds a baseline disparity of 0.92% vs. 0.49%.

Takeaways

Cell-level aggregate consistency completes the calibrated-MRP lineage: state shifts (Ghitza–Gelman 2013) → county (2020) → race × CD (Kuriwaki et al. 2024) → the cell itself.

Race remains essential — but no longer sufficient. The axes that carry much of the identifying variation are within racial groups.

Costs, named: cell-level posteriors are wider by construction; comparative settings need the IPF step and auxiliary priors.

ssrn.com/abstract=6610738