

Monte Carlo Validation of the RSDC Estimators

Parameter recovery under RSDC 1.7.0 (reparameterized global search)

July 18, 2026

1 Introduction

This note validates the RSDC estimators by Monte Carlo parameter recovery: for each model we simulate from a known data-generating process, re-estimate, and check that bias and RMSE shrink as T grows. Five specifications are studied: the constant model, the two- and three-regime fixed-transition (noX) and TVTP models. Relative to the previous validation note, the estimators now use the 1.7-0 *reparameterized global search* (canonical partial correlations; bounded-softmax noX head; top-3 refinement), the case-4 dimension sweep extends to $K = 5$, and the multimodal $N = 3$ cases are estimated under two arms: the default single search and the documented multi-start protocol (`control$n_starts = 4`). The DGPs, true parameters and base seeds are unchanged, so every pre-existing cell re-estimates the *identical simulated data sets* as the previous study.

2 Design and metrics

For every cell we run $M = 100$ replications, each self-seeded, parallelised over 14 cores (452 minutes in total). For an estimate $\hat{\theta}$ of a true value θ we report the mean, bias, RMSE and SD across replications; $\text{RMSE}^2 \approx \text{bias}^2 + \text{SD}^2$, so $\text{RMSE} \approx \text{SD}$ means sampling variance dominates. Consistency is evidenced by bias and RMSE shrinking toward 0 as T grows. TVTP replications whose coefficients reach the ± 10 optimiser bound are excluded from the summaries (see Section 7); regime labels are matched by the ascending-correlation ordering.

3 Case 1: constant correlation ($K = 2$)

Table 1: Case 1 (const): recovery of ρ (true = 0.50).

param	true	mean	bias	RMSE	SD
$T = 500$ (<i>usable</i> = 100)					
ρ	0.500	0.4989	-0.0011	0.0284	0.0285
$T = 1000$ (<i>usable</i> = 100)					
ρ	0.500	0.5032	0.0032	0.0203	0.0202
$T = 2000$ (<i>usable</i> = 100)					
ρ	0.500	0.5012	0.0012	0.0140	0.0140

Essentially unbiased at all T ; RMSE halves as T quadruples, the textbook $\sqrt{T}$ rate. All 100/100 replications converge.

4 Case 2: fixed transitions, $N = 2$

Table 2: Case 2 (noX, $N = 2$): parameter recovery.

param	true	mean	bias	RMSE	SD
$T = 500$ (<i>usable</i> = 100)					
p_{11}	0.900	0.8063	-0.0937	0.2708	0.2554
p_{22}	0.800	0.7359	-0.0641	0.2646	0.2581
ρ_1	0.200	0.1402	-0.0598	0.2283	0.2214
ρ_2	0.700	0.7026	0.0026	0.1353	0.1360
$T = 1000$ (<i>usable</i> = 100)					
p_{11}	0.900	0.8843	-0.0157	0.0905	0.0896
p_{22}	0.800	0.7984	-0.0016	0.1136	0.1141
ρ_1	0.200	0.1699	-0.0301	0.1239	0.1208
ρ_2	0.700	0.7090	0.0090	0.0973	0.0974
$T = 2000$ (<i>usable</i> = 100)					
p_{11}	0.900	0.8832	-0.0168	0.0684	0.0666
p_{22}	0.800	0.7969	-0.0031	0.0877	0.0881
ρ_1	0.200	0.1779	-0.0221	0.0676	0.0642
ρ_2	0.700	0.6970	-0.0030	0.0633	0.0635

Transition probabilities and both regime correlations are recovered with shrinking bias and RMSE; 100/100 replications converge at every T .

5 Case 3: time-varying transitions, $N = 2$

The usable count rises with T ($48 \rightarrow 88 \rightarrow 98$ of 100; the previous study had $46 \rightarrow 86 \rightarrow 98$ on the same data): identification of the covariate slopes sharpens with the sample size, and the reparameterized search slightly increases the usable share at small T .

6 Case 4: fixed transitions, $N = 3$, $K \in \{2, 3, 5\}$

All 100/100 replications converge in every cell, *including* $K = 5$ (33 parameters) — a dimension at which the pre-1.7-0 natural-space search produced no usable cell at all. Median RMSE *falls* with K (Figure 1): the higher-dimensional gains conjectured in the paper are visible in simulation. The multi-start arm (`n_starts` = 4) reproduces these results within Monte Carlo noise: on these well-separated DGPs the default single search already finds the optimum, and the protocol acts as insurance.

7 Case 5: time-varying transitions, $N = 3$

Regime labels are correct in *all* 100/100 replications at both T — no label switching. The usable filter is driven entirely by the ± 10 coefficient bound: 75/100 replications reach it at $T = 1000$, 36/100 at $T = 2000$. This is the classical *separation* problem of multinomial-logit transitions — with persistent regimes some transitions are observed only a handful of times, and their logit MLE diverges — and it fades as T grows. Consistent with that reading, the multi-start arm reaches the

Table 3: Case 3 (tvtp, $N = 2$): recovery on usable replications.

param	true	mean	bias	RMSE	SD
$T = 500$ ($usable = 48$)					
p_{11}	0.700	0.6554	-0.0446	0.2958	0.2955
p_{22}	0.700	0.7028	0.0028	0.2311	0.2336
$\beta_{1,1}$	0.847	0.5728	-0.2745	2.8884	2.9058
$\beta_{1,2}$	-0.800	-1.6250	-0.8250	2.4805	2.3641
$\beta_{2,1}$	0.847	1.3513	0.5040	1.9723	1.9270
$\beta_{2,2}$	0.800	1.6607	0.8607	2.3366	2.1953
ρ_1	0.200	0.1219	-0.0781	0.2308	0.2195
ρ_2	0.700	0.7141	0.0141	0.0898	0.0896
$T = 1000$ ($usable = 88$)					
p_{11}	0.700	0.6475	-0.0525	0.2162	0.2110
p_{22}	0.700	0.6504	-0.0496	0.1995	0.1943
$\beta_{1,1}$	0.847	0.8329	-0.0144	1.3716	1.3794
$\beta_{1,2}$	-0.800	-1.3007	-0.5007	1.7379	1.6737
$\beta_{2,1}$	0.847	0.7379	-0.1094	1.3534	1.3567
$\beta_{2,2}$	0.800	1.0209	0.2209	1.8050	1.8017
ρ_1	0.200	0.1610	-0.0390	0.1383	0.1334
ρ_2	0.700	0.7020	0.0020	0.0675	0.0679
$T = 2000$ ($usable = 98$)					
p_{11}	0.700	0.6884	-0.0116	0.1155	0.1155
p_{22}	0.700	0.7041	0.0041	0.1013	0.1017
$\beta_{1,1}$	0.847	0.8537	0.0064	0.5884	0.5914
$\beta_{1,2}$	-0.800	-0.9666	-0.1666	0.7691	0.7547
$\beta_{2,1}$	0.847	0.9197	0.0724	0.5242	0.5218
$\beta_{2,2}$	0.800	1.0338	0.2338	0.8825	0.8553
ρ_1	0.200	0.1815	-0.0185	0.0794	0.0776
ρ_2	0.700	0.7018	0.0018	0.0384	0.0385




Figure 1: Median RMSE across the eight case-4 parameters, by K and T . Adding series *improves* recovery: each extra series contributes information about the same latent regime path.

Table 4: Case 4 (noX, $N = 3$), $K = 2$: recovery (default search).

param	true	mean	bias	RMSE	SD
$T = 500$ ($usable = 100$)					
p_{11}	0.900	0.5729	-0.3271	0.5148	0.3995
p_{22}	0.850	0.6246	-0.2254	0.4513	0.3930
p_{33}	0.800	0.6698	-0.1302	0.3492	0.3256
ρ_1	0.200	-0.0090	-0.2090	0.4245	0.3713
ρ_2	0.500	0.4581	-0.0419	0.2135	0.2104
ρ_3	0.800	0.8071	0.0071	0.1158	0.1161
$T = 1000$ ($usable = 100$)					
p_{11}	0.900	0.7238	-0.1762	0.3837	0.3426
p_{22}	0.850	0.7368	-0.1132	0.3338	0.3156
p_{33}	0.800	0.6549	-0.1451	0.3287	0.2965
ρ_1	0.200	0.0676	-0.1324	0.3218	0.2947
ρ_2	0.500	0.4860	-0.0140	0.1578	0.1580
ρ_3	0.800	0.8314	0.0314	0.1029	0.0985
$T = 2000$ ($usable = 100$)					
p_{11}	0.900	0.7322	-0.1678	0.3592	0.3192
p_{22}	0.850	0.7548	-0.0952	0.3302	0.3178
p_{33}	0.800	0.7191	-0.0809	0.2819	0.2714
ρ_1	0.200	0.0906	-0.1094	0.2962	0.2766
ρ_2	0.500	0.4793	-0.0207	0.1363	0.1354
ρ_3	0.800	0.7992	-0.0008	0.0777	0.0781

Table 5: Case 4 (noX, $N = 3$), $K = 3$: recovery (default search).

param	true	mean	bias	RMSE	SD
$T = 500$ ($usable = 100$)					
p_{11}	0.900	0.6241	-0.2759	0.4531	0.3613
p_{22}	0.850	0.4924	-0.3576	0.5256	0.3872
p_{33}	0.800	0.7251	-0.0749	0.2848	0.2761
ρ_1	0.200	0.1345	-0.0655	0.2005	0.1904
ρ_2	0.500	0.4585	-0.0415	0.1552	0.1503
ρ_3	0.800	0.7422	-0.0578	0.1046	0.0876
$T = 1000$ ($usable = 100$)					
p_{11}	0.900	0.6835	-0.2165	0.4001	0.3381
p_{22}	0.850	0.5528	-0.2972	0.4726	0.3693
p_{33}	0.800	0.7461	-0.0539	0.2406	0.2356
ρ_1	0.200	0.1593	-0.0407	0.1972	0.1939
ρ_2	0.500	0.4777	-0.0223	0.1489	0.1479
ρ_3	0.800	0.7642	-0.0358	0.0679	0.0579
$T = 2000$ ($usable = 100$)					
p_{11}	0.900	0.7200	-0.1800	0.3664	0.3208
p_{22}	0.850	0.6629	-0.1871	0.3887	0.3424
p_{33}	0.800	0.7628	-0.0372	0.1913	0.1886
ρ_1	0.200	0.1695	-0.0305	0.1820	0.1803
ρ_2	0.500	0.4694	-0.0306	0.1453	0.1427
ρ_3	0.800	0.7802	-0.0198	0.0494	0.0454

Table 6: Case 4 (noX, $N = 3$), $K = 5$: recovery (default search).

param	true	mean	bias	RMSE	SD
$T = 500$ (<i>usable</i> = 100)					
p_{11}	0.900	0.5157	-0.3843	0.5198	0.3517
p_{22}	0.850	0.6436	-0.2064	0.3906	0.3333
p_{33}	0.800	0.8195	0.0195	0.1536	0.1531
ρ_1	0.200	0.1375	-0.0625	0.1552	0.1428
ρ_2	0.500	0.3697	-0.1303	0.1842	0.1308
ρ_3	0.800	0.7297	-0.0703	0.0953	0.0647
$T = 1000$ (<i>usable</i> = 100)					
p_{11}	0.900	0.6604	-0.2396	0.4093	0.3335
p_{22}	0.850	0.5857	-0.2643	0.4464	0.3615
p_{33}	0.800	0.8133	0.0133	0.1074	0.1071
ρ_1	0.200	0.1857	-0.0143	0.1223	0.1221
ρ_2	0.500	0.4116	-0.0884	0.1475	0.1186
ρ_3	0.800	0.7567	-0.0433	0.0622	0.0448
$T = 2000$ (<i>usable</i> = 100)					
p_{11}	0.900	0.7383	-0.1617	0.3430	0.3040
p_{22}	0.850	0.5902	-0.2598	0.4515	0.3711
p_{33}	0.800	0.8119	0.0119	0.0784	0.0779
ρ_1	0.200	0.1994	-0.0006	0.1041	0.1047
ρ_2	0.500	0.4417	-0.0583	0.1242	0.1101
ρ_3	0.800	0.7657	-0.0343	0.0492	0.0354

bound slightly *more* often (84 and 43 of 100): a better search locates the boundary MLE more reliably, so the bound filter screens difficult *samples*, not optimizer failures.

8 Comparison with the pre-1.7-0 study

Same DGPs, same seeds, hence identical simulated data; only the estimator changed. Cases 1–2 are numerically indistinguishable (unimodal surfaces). Case 3 gains usable replications at small T (48 vs 46 at $T = 500$; 88 vs 86 at $T = 1000$) with comparable RMSE. Case 5 usable rates are essentially unchanged (25/64 vs 27/64 of 100), as expected since the bound filter reflects the data, not the search. The decisive differences are structural: the $K = 5$ cell of case 4 (infeasible before) now converges 100/100 with the smallest RMSEs of the whole sweep, and the full study completes in 452 minutes on 14 cores.

9 Conclusion

The 1.7-0 estimators recover the true parameters of all five specifications, with textbook $\sqrt{T}$ behaviour where surfaces are regular, monotone gains in T elsewhere, and *improving* accuracy in the cross-section dimension. The reparameterized search removes the high-dimensional feasibility barrier ($K = 5$, 33 parameters, 100/100) while leaving well-identified low-dimensional cells untouched. Remaining finite-sample limitations (coefficient-bound separation in sparse-transition TVTP cells) are data features that vanish with the sample size and are flagged, not hidden, by the reported usable counts.

Table 7: Case 5 (tvtp, $N = 3$, $K = 2$): recovery on usable replications (default search).

param	true	mean	bias	RMSE	SD
<i>T</i> = 1000 (<i>usable</i> = 25)					
p_{11}	0.900	0.8564	-0.0436	0.1921	0.1910
p_{22}	0.900	0.8409	-0.0591	0.2281	0.2248
p_{33}	0.900	0.8907	-0.0093	0.0495	0.0496
$\beta_{1,1}$	3.401	4.7837	1.3825	2.8500	2.5436
$\beta_{1,2}$	0.500	0.7728	0.2728	1.3825	1.3832
$\beta_{1,3}$	0.847	2.2351	1.3878	3.5132	3.2940
$\beta_{1,4}$	0.000	0.1243	0.1243	2.8442	2.9001
$\beta_{2,1}$	0.000	-0.9922	-0.9922	4.0514	4.0090
$\beta_{2,2}$	0.000	-0.5460	-0.5460	2.7040	2.7029
$\beta_{2,3}$	2.890	3.2251	0.3347	1.8514	1.8585
$\beta_{2,4}$	0.500	0.9401	0.4401	2.5719	2.5862
$\beta_{3,1}$	-3.401	-4.1285	-0.7273	2.0801	1.9890
$\beta_{3,2}$	-0.500	-0.1897	0.3103	1.3535	1.3446
$\beta_{3,3}$	-2.554	-2.9883	-0.4344	1.6455	1.6199
$\beta_{3,4}$	0.000	-0.1678	-0.1678	1.4904	1.5115
ρ_1	-0.600	-0.5954	0.0046	0.0835	0.0851
ρ_2	0.100	0.1396	0.0396	0.1448	0.1421
ρ_3	0.800	0.8091	0.0091	0.0425	0.0424
<i>T</i> = 2000 (<i>usable</i> = 64)					
p_{11}	0.900	0.8991	-0.0009	0.1059	0.1067
p_{22}	0.900	0.8913	-0.0087	0.1050	0.1055
p_{33}	0.900	0.9003	0.0003	0.0247	0.0249
$\beta_{1,1}$	3.401	3.8880	0.4868	1.4910	1.4205
$\beta_{1,2}$	0.500	0.4291	-0.0709	0.9553	0.9602
$\beta_{1,3}$	0.847	0.3253	-0.5220	3.1084	3.0884
$\beta_{1,4}$	0.000	-0.0299	-0.0299	2.4781	2.4975
$\beta_{2,1}$	0.000	-0.1718	-0.1718	1.9028	1.9100
$\beta_{2,2}$	0.000	0.0826	0.0826	1.0031	1.0076
$\beta_{2,3}$	2.890	3.1477	0.2573	0.9994	0.9733
$\beta_{2,4}$	0.500	0.7555	0.2555	1.1294	1.1088
$\beta_{3,1}$	-3.401	-4.1598	-0.7586	1.6758	1.5061
$\beta_{3,2}$	-0.500	-0.4501	0.0499	1.1058	1.1134
$\beta_{3,3}$	-2.554	-2.6610	-0.1071	0.6608	0.6573
$\beta_{3,4}$	0.000	0.0368	0.0368	0.6354	0.6393
ρ_1	-0.600	-0.5958	0.0042	0.0724	0.0728
ρ_2	0.100	0.1151	0.0151	0.0962	0.0958
ρ_3	0.800	0.8015	0.0015	0.0196	0.0197