

From M Passes to One: A MapReduce-Style Bootstrap for Pointwise Off-Policy Evaluation in Production Recommenders

Ghulam Ahmed Ansari and Cheng Ding, LinkedIn Corporation. RecSys 2026 Industry Track.

[Download the print poster PDF \(24 x 36 inches\).](#)

[Download the text-and-data PDF \(tagged\).](#)

This page is the text alternative to the visual poster. The tagged PDF presents the same explanation and tables in document form; it is separate from the large-format print poster.

- [Production logging](#)
- [One-pass bootstrap](#)
- [Pointwise and cascade comparison](#)
- [Score calibration](#)
- [Online agreement](#)

From dedicated random sessions to production logs

The deployed system combines quantile calibration, pointwise IIPS@K, and one-pass uncertainty estimation. It uses full-slate production logs, enabling retirement of the dedicated random-session branch. Position-one replay baseline: uniform-random logging (WSDM 2011), evaluated at the first position. Requires logging support and propensity assumptions.

Make uncertainty practical: 50 passes become one

M separate reads are replaced by one shared read. The illustration shows tile 1, tile 2, and tile M: same rows in every tile, with a different seeded binary mask. $M = 50$ replicas. Columns are mask slots, not slate positions. Mask slots: 0 1 0 1 $\rightarrow$ session weight 2 (illustrative). The illustrated masks are conceptual, not experiment data.

1. Read each session once and compute $g(s)$, the session's OPE contribution.
2. Fan that contribution into all tiles. Sum each row's bits for its weight, then multiply $g(s)$ by that weight.
3. Keep replicas separate: the 50 masked accumulators collect weighted contributions separately by replica.
4. Normalize each sum by the common mask-slot count times session-count denominator. This produces M estimates.
5. Compute sample variance across the estimates.

The collection line routes separate replicas; it does not pool them into one estimate. Masks and updates still scale with M.

Within 1% relative variance error vs. same-mask multi-pass reference. $M = 50$ is the number of replicas; K is slate depth. The factorization reduces I/O passes, not a measured 50x wall-clock speedup. Detailed replica-count and dropout sweeps are outside this poster's scope; this comparison does not establish formal confidence-interval coverage.

Low pointwise error through $K = 60$

Under the pointwise assumption, reward at a position depends on its item without modeling dependence on preceding items. Pointwise IPS is established prior work; the empirical comparison tests its suitability for this Feed workload.

Pointwise has lower error here; cause not isolated. Unstable sequential weights are one possible reason. Higher nMSE alone does not distinguish weight instability, propensity error, and a misspecified dependence model, or prove that cross-position reward effects are absent.

Normalized mean squared error: Feed composite engagement against independent on-policy logs; lower is better.

Slate depth K	Pointwise IIPS@K	Cascade RI-IPS@K
1	1.8e-6	2.1e-6
10	2.0e-6	3.4e-5
20	4.1e-6	1.7e-4
60	6.8e-6	2.9e-3

These values are from the accepted paper's composite-engagement ablation table. The poster uses a logarithmic vertical axis, linear K spacing, and distinct colors, marker shapes, and line styles. Pointwise error stays low through K = 60; the tested cascade estimator deteriorates. This is a workload-specific finding, not a universal ranking of assumptions.

Calibrate scores before computing propensities

Across ranker versions, objective re-weights can change score scales. Raw model scores (logits) passed directly to softmax can confuse scale changes with preference changes. The calibrated path is: equal-population score bins (quantile bins with equal item counts), observed response rates, then softmax. The resulting propensity estimates are used for weighting. Calibrate candidate and logging rankers separately. Weight: candidate / logging probability (π_e / π_l). Softmax is retained.

The schematic is conceptual; it contains no measured score distributions. Rank-preserving re-scaling is the lemma's guarantee. Arbitrary re-weights may change order; exact raw-to-calibrated AUC preservation is not asserted.

Section 6.4: agreement with online direction in two comparisons across three ranker generations.

Comparison	Raw-score policy	Calibrated policy
Within-regime	Matches	Matches
Across re-weight	Differs	Matches

Online-direction agreement: $1/2 \rightarrow 2/2$. The within-regime comparison is M2 versus M1; the across-reweight comparison is M3 versus M2. This small ablation is separate from the promotion-parity cohorts below.

Does the complete system agree with online outcomes?

Offline-online agreement (called parity in the paper): matching the online direction and significance verdict.

Method / surface	Matches	Rate	Wilson 95% interval
Position-one replay baseline	2/11	18%	5-48%
Bootstrap IIPS@K	7/11	64%	35-85%
Suggested Post	9/12	75%	47-91%

The first two rows compare the same 11 second-pass pairs: 64% (7/11) versus 18% (2/11). Suggested Post is a separate 12-comparison model set. Combined offline-online agreement (two surfaces): 70% (16/23), Wilson 95% CI 49-84%. Kendall

rank correlation (τ) is +0.29 ($p=0.04$) in a broader set of 26 model comparisons with both estimates populated.

Comparison scope

The replay baseline and proposed system use different logging data and slate depths. End-to-end agreement reflects the complete system; individual contributions to that result were not isolated.

Candidate-selection boundary. The estimator does not correct bias from earlier candidate selection. It does not recover responses for candidates excluded upstream. This limits what changes can be evaluated from the available logs; it does not establish that estimates within the evaluated scope are wrong.

Presentation and citation

P1, Tuesday September 29, 2026, Atrium / Trees. Presentation windows: 10:00-10:45 AM and 3:30-4:00 PM, Minneapolis local time. Set up 9:00-10:00 AM; remove at 4:00 PM.

[Official conference listing](#)

Paper DOI: 10.1145/3773078.3831911.