Efficient CVaR Estimation for Mortgage Pools via Static Importance Sampling

Olateju Ojeyinka: Bloomberg LP, MS Candidate [Financial Engineering / Columbia University]
Contact: oao2134@Columbia.edu GitHub: github.com/ojeyio/mbs_pools_cvar_is

Introduction / Motivation

- Accurate CVaR estimation is essential for risk management of MBS pools.
- Naive MC requires large draws to resolve α = 0.99 tails which is computationally expensive.
- Static IS (cheap exponential tilt) can reduce variance and increase effective tail sampling while remaining simple to deploy on loan-level cashflows.

Data

Simulated loan-level cohort designed to match 2015–2019 origination characteristics: 30-yr fixed loans, realistic UPB distribution, coupon, LTV, and FICO. Cohort size used in experiments: **N_loans = 300**. The pipeline accepts a standard loan CSV with these required columns: loan_id, origination_year, original_term_months, original_upb, current_upb, note_rate, orig_ltv, fico.

Methods

- Compute per-loan PV cashflows with PSA-style prepayment and stressed market rates.
- Portfolio loss L(z) = PV_base − PV_stressed for shock z.
- IS sample z ∼ N(θ,1), weight by likelihood ratio. Calibrate θ by pilot runs to reduce CVaR variance.

Pseudo-code / Key equations

Sampling & Weighting (exponential tilt)

- Sample: $z \sim \mathcal{N}(\theta, 1)$.
- Importance weight: $w(z) = \frac{\phi(z)}{\phi(z-\theta)} = \exp\left(-\theta z + \frac{1}{2}\theta^2\right)$.
- Estimate CVaR: use weighted tail estimator at level α with normalized weights.

Algorithm (high-level)

1. Pilot: draw z_i (i=1..pilot_N), compute losses L(z_i), choose θ candidate.
2. Sensitivity sweep: run weight diagnostics (ESS, ESS_tail, VaR_w, CVaR_w) for θ grid.
3. Final: run experiments with chosen θ for repeats & bootstrap.

Calibration & Diagnostics

- Pilot calibration: θ was pilot-calibrated by running pilot_N draws and selecting θ that maximized tail information while keeping overall ESS stable. Sensitivity sweep performed over θ ∈ {0.4, 0.5, 0.6}.
- Example diagnostics (pilot run shown):
  - Pilot θ ≈ **0.5**: **ESS = 7,754** (ESS/N ≈ 0.775), **Tail ESS = 331**, Naïve tail count = 100.
  - Pilot CVaR (α=0.99): **$31,919**.
  - Bootstrap ΔCVaR (IS − MC) mean ≈ **$317,184**, 95% CI = [$126,726, $518,949].

Note: *sensitivity results show ϑ = 0.6 can further increase Tail ESS; final poster will use ϑ =* ***0.6***

θ	ESS (total)	ESS / N	Tail ESS	Naive tail count	VaR_w	CVaR_w
0.4	8,504	0.850	266.6	100	28,534	31,910
0.5	7,754	0.775	330.5	100	28,474	31,919
0.6	6,920	0.692	407.1	100	28,560	31,918

Calibration sweep (N = 10,000 per experiment, α = 0.99). Increasing ϑ moves sampling mass into the tail: Tail ESS rises from 266 → 331 → 407 (ϑ = 0.4 → 0.6). Overall ESS declines with larger ϑ but remains healthy (ESS/N ≥ 0.69). VaR and CVaR point estimates remain stable across the sweep.

Experimental Setup

- Cohort: simulated 2015–2019 loans (N_loans = 300).
- CVaR level: α = 0.99.
- Monte Carlo draws per experiment: N = 10,000.
- Pilot-calibrated θ: ≈ 0.5 (sensitivity shows θ up to 0.6 increases tail ESS).
- Repeats and pilot sizes chosen to provide bootstrap CI for mean difference.

Key Results

- Variance Reduction Ratio (VRR) ≈ 6× (examples: ~5.98 at N=5k; ~6.8 at N=10k).
- Final θ ≈ 0.6: ESS ≈ 6,920 (ESS/N ≈ 0.692) and tail ESS ≈ 407 (naive tail count = 100).
- Weighted CVaR (α=0.99) ≈ $31,919
- Bootstrap mean difference (IS − MC) in CVaR = **$317,183.77**, 95% CI = **[$126,726.02, $518,948.71]**.
- Interpretation: IS both reduces estimator variance and reveals additional tail mass that naive MC under-sampled.

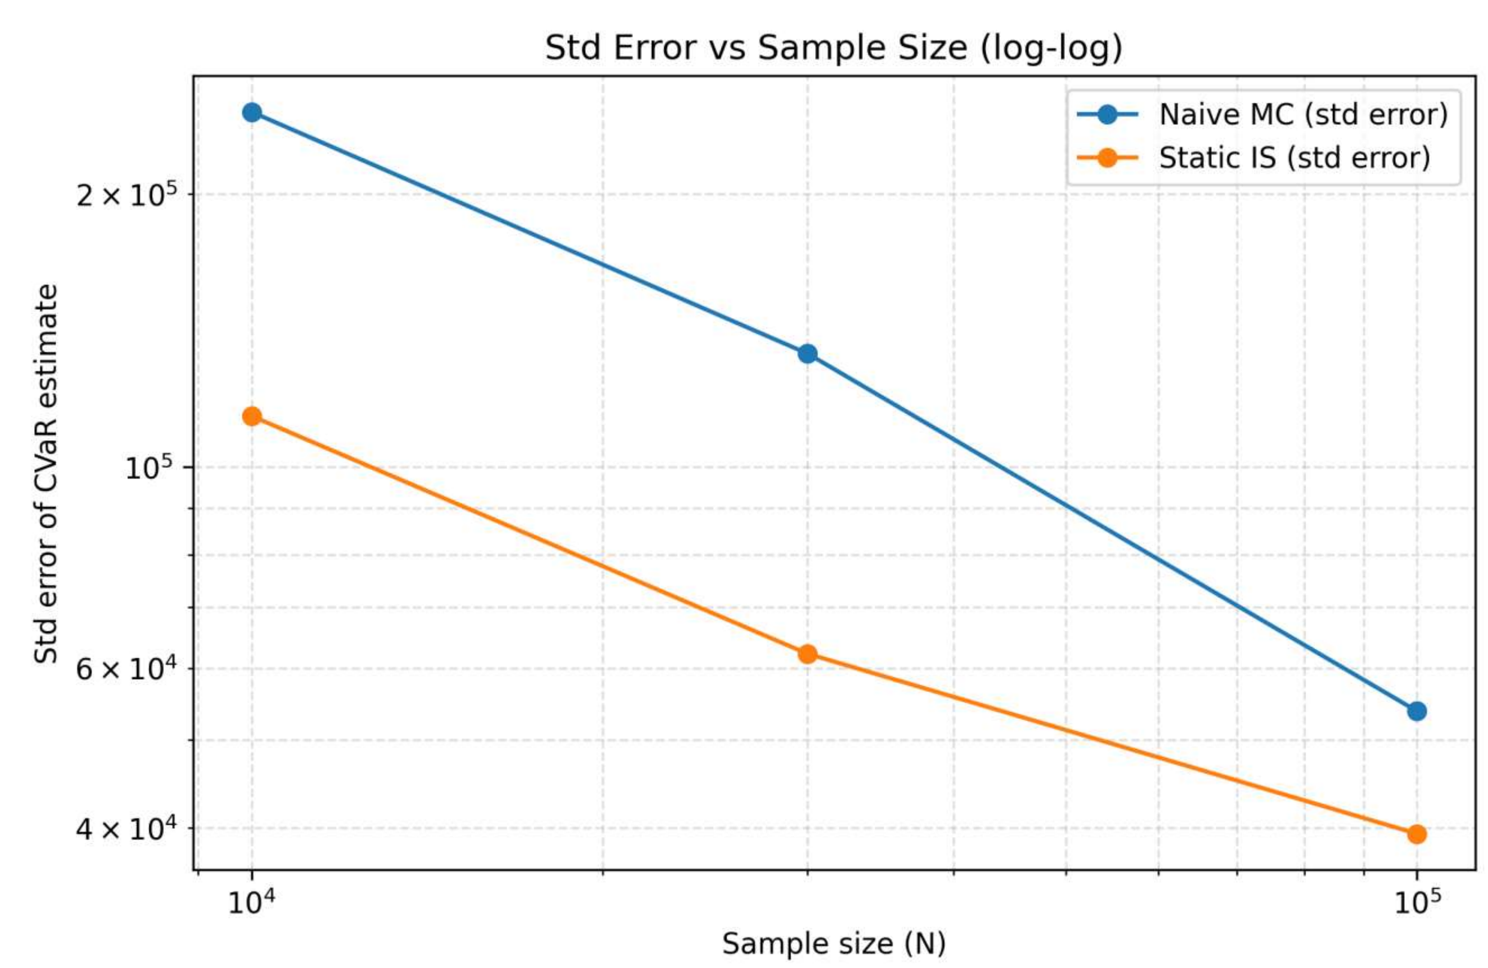

Figure 1 (Variance plot): *Standard error of estimated $CVaR_{0.99}$ vs Monte Carlo sample size (log–log). Static IS (final ϑ ≈ 0.6) reduces variance by ≈5x compared to naive MC*

Run (θ)	Repeats	IS mean (CVaR)	IS std	MC mean (CVaR)	MC std	Mean diff (IS − MC)	Rel diff (%)	VRR ≈ Var(MC)/Var(IS)
pilot (0.5)	30	16,615,472	136,724	16,424,067	356,781	**191,405**	**+1.16 %**	**≈ 6.81×**
final (0.6)	90	16,566,462	77,943	16,556,737	168,959	**9,725**	**+0.06 %**	**≈ 4.70×**

Comparison of $CVaR_{0.99}$ estimates (IS vs naive MC). Pilot runs (θ ≈ 0.5) showed IS had higher mean and much lower std (variance reduction ≈ 6.8×). Final validated runs (θ = 0.6, more repeats) show the point estimates converge (mean diff ≈ $9.7k, practically negligible) while IS maintains a large variance reduction (≈ 4.7×). These diagnostics support choosing θ = 0.6 for stronger tail resolution while preserving estimator stability.

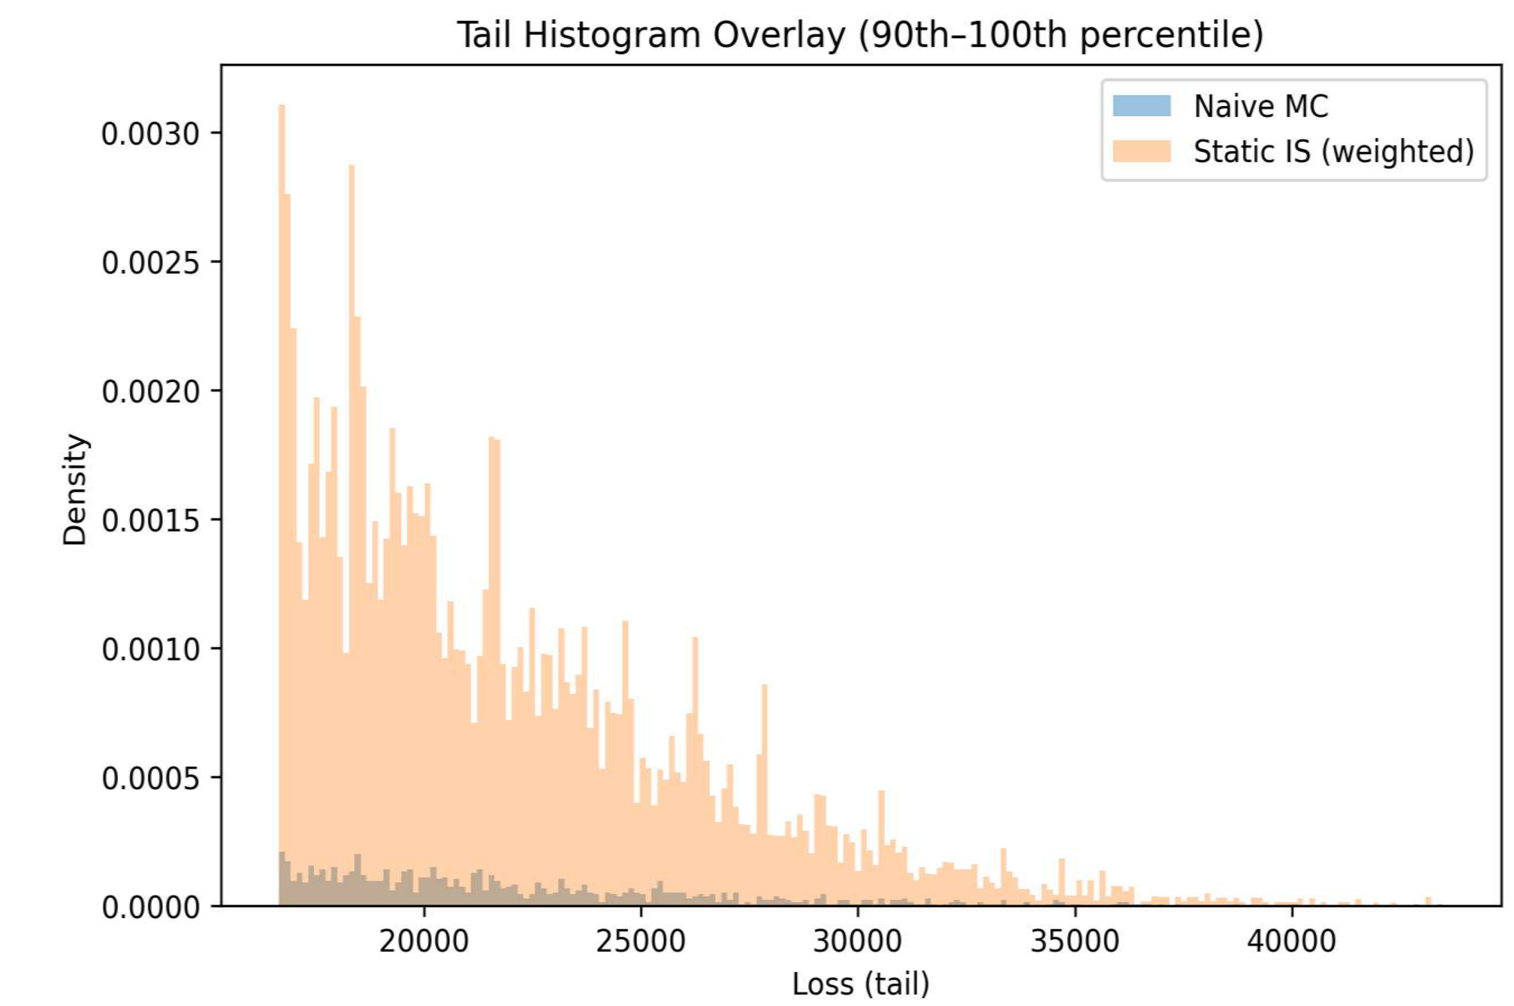

Figure 2 (Tail histogram overlay): *Tail density (90th–100th percentile): naive MC (unweighted) vs IS-weighted density. IS increases sampling density in extreme losses — weights correct for bias and improve tail resolution.*

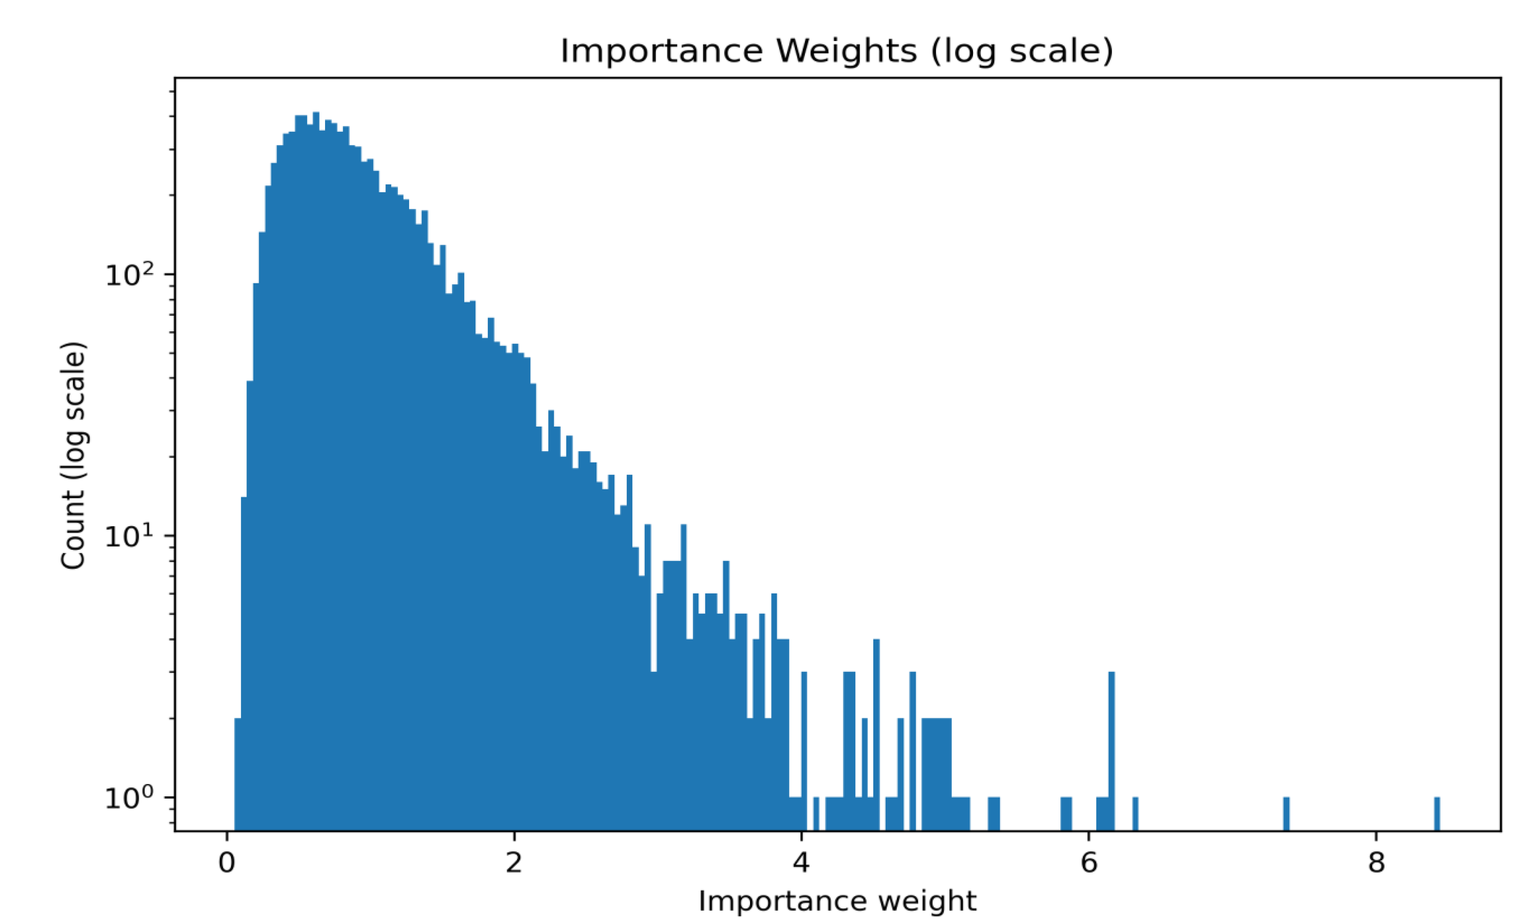

Figure 3 (Importance weights, log scale): *Distribution of importance weights (log scale). We observe no single dominating weight; overall ESS ≈ 7,754 and tail ESS ≈ 331 for the pilot ϑ ≈ 0.6 run.*

Conclusions

- Static IS with a calibrated exponential tilt gives large variance reductions for CVaR in MBS pools.
- IS substantially increases effective tail sample size (ESS_tail ≫ naive tail count), improving tail estimation quality.
- Approach is compatible with loan-level CSVs and production cashflow engines — ready to scale to GSE datasets.

Future Work

- Run pipeline on full Freddie/Fannie acquisition + performance files.
- Explore adaptive / mixture proposals and learned proposals (GAN-based) to further increase efficiency.
- Refine prepayment & interest-rate dynamics rather than using a simple PSA-style CPR mapping.