

The Layla-Sphere Method: Dependent Rounding for the Landing Phase of Balanced Random Sampling

Abdelkader Bousabaa¹ & Rémi Sirolli²

¹ENSAI – University of Paris-Saclay

²ENSAI – Crédit Agricole CIB, Head of Strategic Management Control

June 16, 2026

Abstract

The Cube method (Deville and Tillé, 2004) achieves balanced random sampling by combining a martingale flight phase with a greedy landing phase that rounds the (at most p) fractional inclusion probabilities remaining after the flight, independently and one at a time. This landing is the principal source of residual balancing error in practical applications.

We first prove a structural result: any landing that rounds the fractional units independently, with their fractional values as probabilities, induces the same sampling design *whatever the processing order* — greedy, balance-informed or trajectory-informed. The landing cannot be improved by reordering it; only *dependent* rounding can carry information.

This article introduces the **Layla Sphere method**, in which inclusion probabilities are embedded as squared moduli of complex amplitudes on the sphere $\mathcal{S}_{n,N} \subset \mathbb{C}^N$. During the flight, each unit accumulates a *phase* tracking its angular trajectory, and pairs of units accumulate a *coherence* recording how systematically the flight traded one against the other. The landing then resolves the fractional set by pivotal martingale transfers between pairs of units, selected either by a balance-minimising rule or by a phase-only rule that couples the most anti-coherent pairs first. For any selection rule, the Horvitz–Thompson estimator remains unbiased; when the inclusion probabilities are among the balancing variables, the realised sample size equals n almost surely — a property the greedy landing destroys. An exact decomposition separates the two channels through which the new landing affects the estimator variance.

A factorial simulation study (3 240 paired comparisons on shared flights; $N \in \{50, 100, 200\}$, $p \in \{2, 5, 10\}$, three population types) shows that the balance rule reduces the median standardised imbalance by 13% for $p = 10$ (sign test $z = +5.5$), restores the exact sample size in 100% of draws (against 40%), and reduces the root mean squared error of the Horvitz–Thompson estimator by 11% on average and up to 19% on fixed populations. The exact enumerative landing of [Deville and Tillé \(2004\)](#) — a small linear programme, now solvable in milliseconds for moderate p — serves as the ceiling: it reaches 45% of imbalance reduction at $p = 10$, of which the streaming balance rule captures about 29%. A mechanism study shows that uninformed dependent rounding does not improve balancing, that the coherence rule alone recovers about half of the balance-rule gain, and that phase information adds nothing measurable on top of the balance rule — a quantified, and we believe honest, assessment of what trajectory information is worth at the landing.

Keywords: balanced sampling, Cube method, Layla Sphere method, landing phase, dependent rounding, pivotal method, phase coherence, Horvitz–Thompson estimator, simulation study.

MSC 2020: 62D05 (Sampling theory), 62D10 (Missing data; incomplete data), 81P68 (Quantum computation).

JEL: C83 (Survey Methods), C15 (Statistical Simulation).

6 Conclusion

Twenty-five years after its first implementation, the CUBE method remains the reference algorithm for balanced random sampling. The present article proposes a targeted improvement to its most vulnerable component: the landing phase.

The route to that improvement passes through a negative result. As long as the residual fractional units are rounded independently, with their fractional values as probabilities, no processing order — greedy, balance-informed or trajectory-informed — can alter the sampling design (Proposition 3.1). The landing cannot be repaired by reordering it; it must be repaired by making its decisions dependent. The LAYLA-SPHERE method does so with pivotal martingale transfers between pairs of units, and concentrates all the available information — the balancing matrix, or the trajectory coherence accumulated on the complex sphere during the flight — in the one choice that matters: which pair to couple. Unbiasedness of the Horvitz–Thompson estimator is preserved for any pair-selection rule, as proved in Appendix A and confirmed by simulation; when the inclusion probabilities are balanced, the realised sample size equals its target almost surely, where the greedy landing misses it in most draws. In the factorial study, the balance rule reduces the median standardised imbalance by 13% for $p = 10$ balancing variables, and the restoration of the exact sample size reduces the root mean squared error of the Horvitz–Thompson estimator by about 11% on average; the coherence rule, using only phase information accumulated at no cost during the flight, recovers about half of the balancing gain without consulting the balancing data at the landing.

To Joseph Lieber: this paper is, in part, a continuation of what we began together in 1999. The CUBE was a remarkable tool then. It remains one today, and perhaps a little more so.