Conventional solution and proof

Under independent sampling, $\mathbb{E}z_m = 0$ and $\operatorname{Var} z_m = 1/m$. Thus for every fixed m,

$$\mathbb{P}(|z_m| \geq \varepsilon) \leq 1/(m\varepsilon^2).$$

No independence between different z_m is needed to apply this inequality separately. However, for $m \leq n$ the means share their first m observations, and

$$\operatorname{Cov}(z_m, z_n) = \frac{1}{mn} \sum_{i \leq m, j \leq n} \operatorname{Cov}(X_i, X_j) = \frac{1}{n}.$$

They are not independent, since their covariance is nonzero. Their tail-event probabilities cannot therefore be multiplied on an independence assumption. A union bound over a finite set of indices is valid, though often loose. If only the marginal mean and variance are assumed, even $\operatorname{Var} z_m = 1/m$ need not hold: taking all X_i equal gives variance one.

Structural or analytic derivation

The vector of partial means is a triangular linear transform of the samples. Its covariance matrix has entries $1/\max(m, n)$. This operator representation exposes every dependence simultaneously and supplies a direct check on any joint concentration calculation. Chebyshev is a marginal second-moment result and remains valid for each coordinate.

Comparative cost and preferred method

The marginal bound takes constant arithmetic per chosen m. Constructing the full $N \times N$ covariance matrix requires $\Theta(N^2)$ outputs; a lazy formula represents every entry in $O(1)$ storage. The main improvement is correction of a false independence inference, not a runtime claim.