The probability generating function is $(1-p+pz)^n$. Its first two derivatives at $z=1$ give $\mathbb{E}S_n = np$ and $\mathbb{E}[S_n(S_n-1)] = n(n-1)p^2$, hence the same variance. The analytic characteristic-function expansion explains the $n^{-1/2}$ standard-error scale and the normal limiting shape. The continuity theorem is used only for the stated distributional interpretation; the exact variance and finite-sample bound above have complete elementary proofs.

Comparative cost and preferred method

Enumerating all 2^n binary sequences is unnecessary. Independence or two generating-function derivatives give the variance in constant symbolic work and $O(1)$ numerical arithmetic after n, p are supplied. Computing the observed fraction still reads n observations, so it costs $\Theta(n)$. Conventional variance additivity already attains the efficient formula; the transform supplies a unified derivation, not a new speedup.