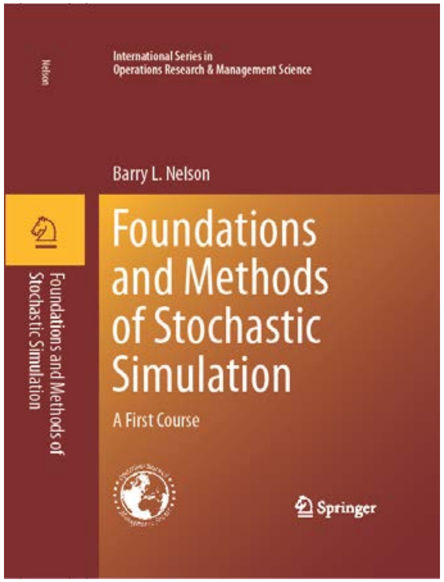

Chapter 5.2:
Simulation as a Stochastic Process

©Barry L. Nelson
Northwestern University

July 2017

Simulated asymptotics

- Good news: Simulation problems can often be treated with "large sample" statistical reasoning because data are cheap and we want high precision.
- Bad news: We don't always have the "i.i.d." output data to which our standard asymptotics apply.
- Here we review/introduce some important ideas that run throughout the remainder of the book.

Thought experiment…

- Suppose each of us is simulating the *M/G/1* queue for one replication and reporting the average waiting time.
- We each independently and randomly select our starting random number seed.
  - Assume a *perfect* generator with infinite seeds.
- We report our averages at four different numbers of customers: m = 10, 100, 1000 and 10,000.
- If we collected the results from all of us (plus enough other people to total 1000), what would we see?

Specifics

We each independently simulate

$$Y_i = \max\{0, Y_{i-1} + X_{i-1} - A_i\}$$

with $Y_0 = X_0 = 0$, and we report

$$\bar{Y}(m) = \frac{1}{m}\sum_{i=1}^{m} Y_i$$

at $m = 10, 100, 1000$ and $10{,}000$.

The parameters I have chosen imply that the steady-state mean is $\mu = 2.133$.

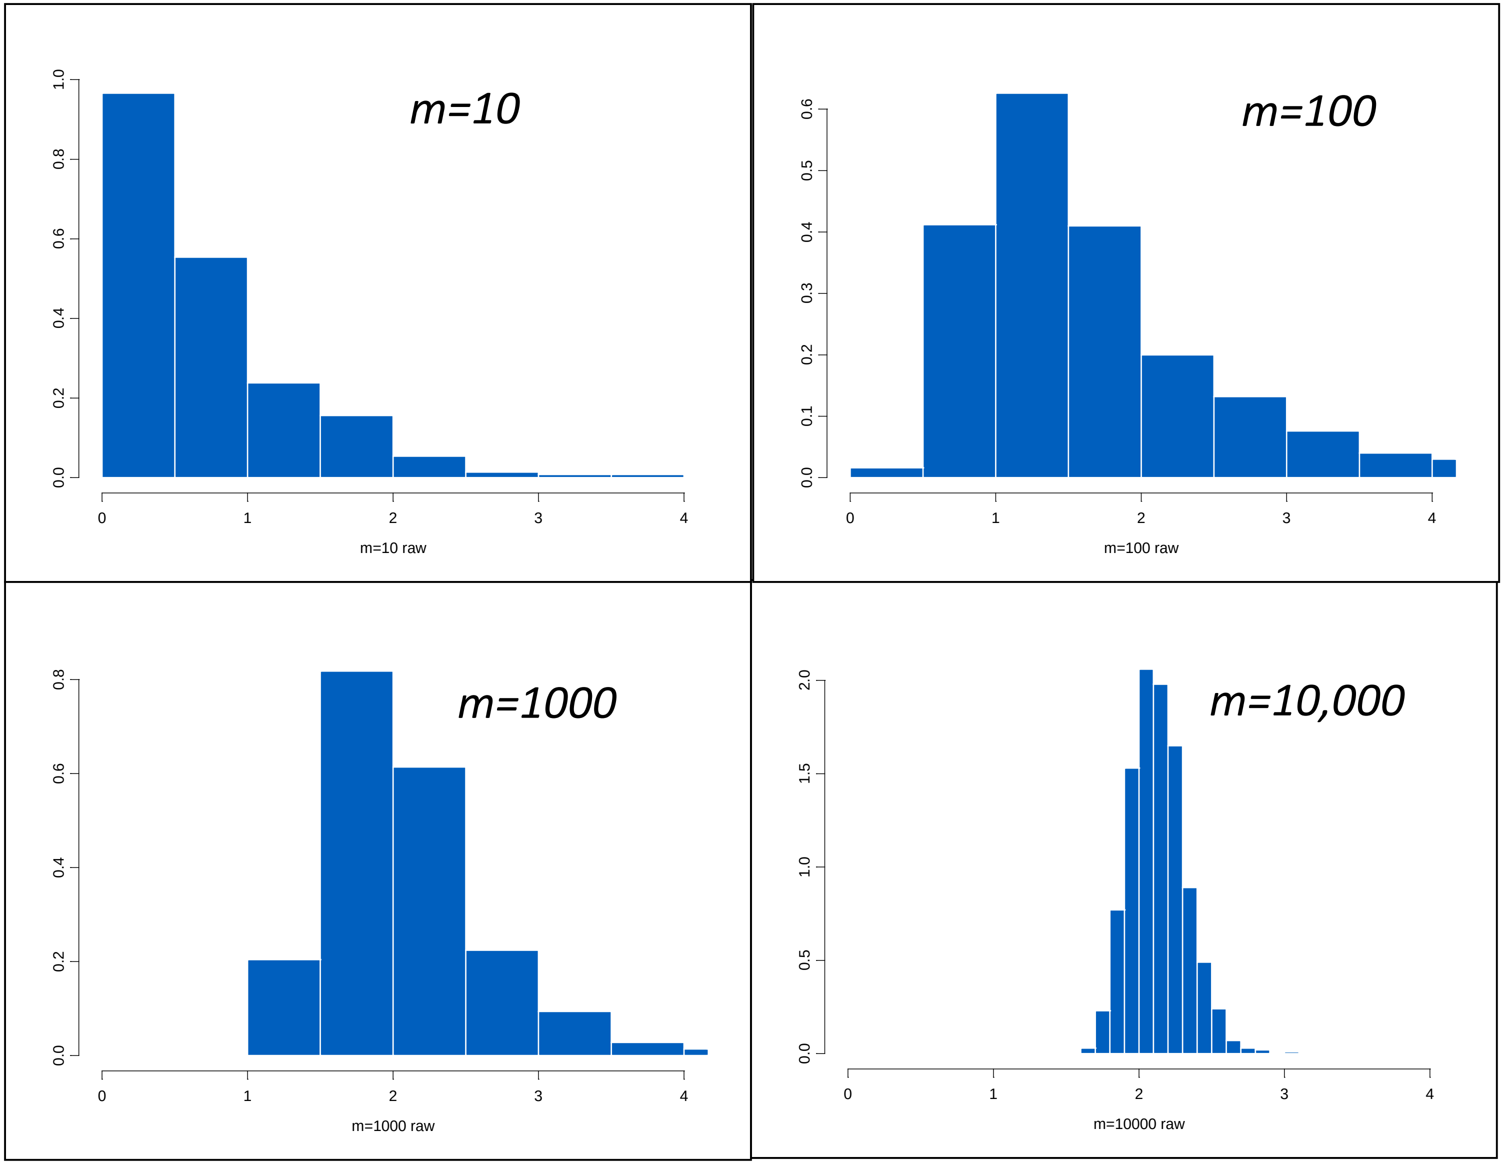

Notice that the 1000 averages are looking more and more normally distributed, but also converging to a single value.

Here we are plotting 1000 values of $\sqrt{m}\left(\overline{Y}(m) - 2.133\right)$

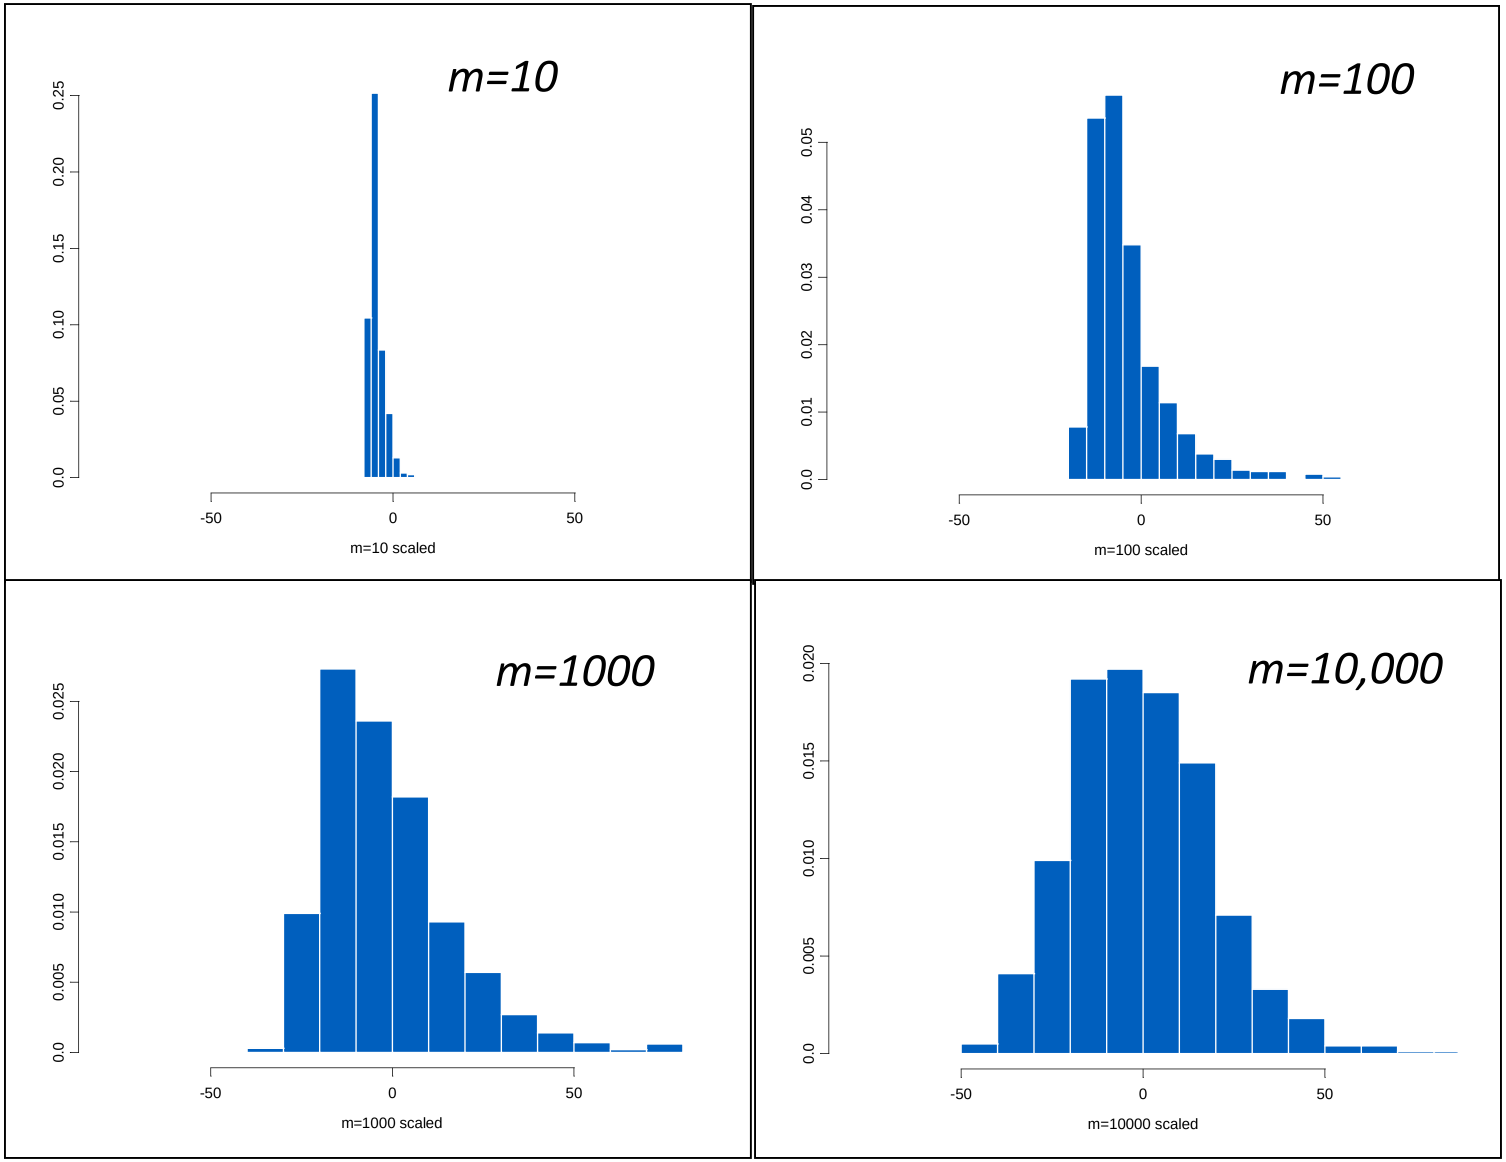

Notice that the scaled and centered averages are looking more and more like a mean 0 stable normal distribution, not converging to a single value.

Summaries of the 1000 averages

1000 averages

0.301432	1.316943	2.181791	2.12834	
0.951291	1.442282	1.60648	2.197148	
0.020483	0.612363	1.68078	2.326293	
0.978482	1.62803	1.738985	2.430449	
1.405882	1.406136	1.362369	2.087307	
0.289488	3.895345	2.148725	2.234362	
0.845869	0.841387	2.062226	2.049821	
0.212686	0.544943	2.001026	2.172981	
0.107063	2.15477	1.470807	2.067342	
1.392107	2.882288	1.67284	2.204039	
1.258184	2.791777	1.684099	1.959779	
1.078646	4.330491	2.746	2.065909	
0.55582	1.02201	1.957363	2.099637	
0.28467	5.880592	2.241638	2.280259	
0.717877	1.738929	2.082655	2.129813	**average**
-1.41512	-0.39407	-0.05034	-0.00319	**bias = average - 2.133**
0.914	0.742	0.359	0.01	**probability > 0.5 off**
m=10	**m=100**	**m=1000**	**m=10000**	

Convergence

- This example illustrates various modes of convergence:
  - Convergence in distribution
    - Scaled average converges to a stable normal distribution.
  - Convergence in probability
    - Chance the average is far away from the mean goes to 0.
  - Convergence with probability 1 (“almost sure convergence”)
    - For "almost all" random number seeds the average converges to the mean.
- But your data were NOT i.i.d.
- We need to understand these modes and when we can expect them to hold in our simulations.

Familiar stuff: Across-rep asymptotics

Strong law of large numbers: If $Z_1, Z_2, \ldots$ are i.i.d. with $\mathrm{E}(Z_1) < \infty$, then

$$\frac{1}{n}\sum_{i=1}^{n} Z_i \xrightarrow{a.s.} \mathrm{E}(Z_1).$$

Central limit theorem: If $Z_1, Z_2, \ldots$ are i.i.d. with $\mathrm{E}(Z_1^2) < \infty$, then

$$\sqrt{n}\left(\frac{1}{n}\sum_{i=1}^{n} Z_i - \mathrm{E}(Z_1)\right) \xrightarrow{D} \gamma\,\mathrm{N}(0,1)$$

where $\gamma^2 = \mathrm{Var}(Z_1)$.

Is E(Z_1) what we want?

SAN: $Z_1 = I(Y_1 > t_p)$ when we want to estimate $\Pr\{Y > t_p\}$? OK

M/G/1:

$$Z_1 = \frac{1}{m} \sum_{j=1}^{m} Y_{1j}$$

when we want to estimate the steady-state mean μ? SLLN and CLT still apply across reps, but convergence is to the wrong value.

This leads to thinking about within-replication asymptotics, as suggested by our group experiment.

Within-replication asymptotics

When might we expect a SLLN or CLT to apply within a single (long) replication? Roughly,...

- Neither the logic nor the input processes change over time, and no periodic (cyclic) behavior.
- The state space of the simulation does not decompose into distinct subsets.
- The state of the simulation in the distant future is effectively independent of the state in the past.

Examples: Stable Markovian queues; irreducible, aperiodic, positive recurrent Markov chains; time series processes; and regenerative processes.

Regenerative processes & steady state

Let Y_t be the output process ($t = 0, 1, 2, \ldots$ or $t \geq 0$).

Suppose in the same simulation we can identify a renewal process $\{S_i, i = 0, 1, 2, \ldots\}$:

$$
\begin{aligned}
S_0 &= 0 \\
S_i &= A_1 + A_2 + \cdots + A_i
\end{aligned}
$$

where A_i are i.i.d. with $\Pr\{A_i = 0\} < 1$ and $\Pr\{A_i < \infty\} = 1$.

The output $\{Y_t\}$ is a *regenerative process* if

$$\{Y_t; S_i \leq t < S_{i+1}\}, i = 0, 1, 2, \ldots$$

are i.i.d.

Regeneration in GI/G/1 queue

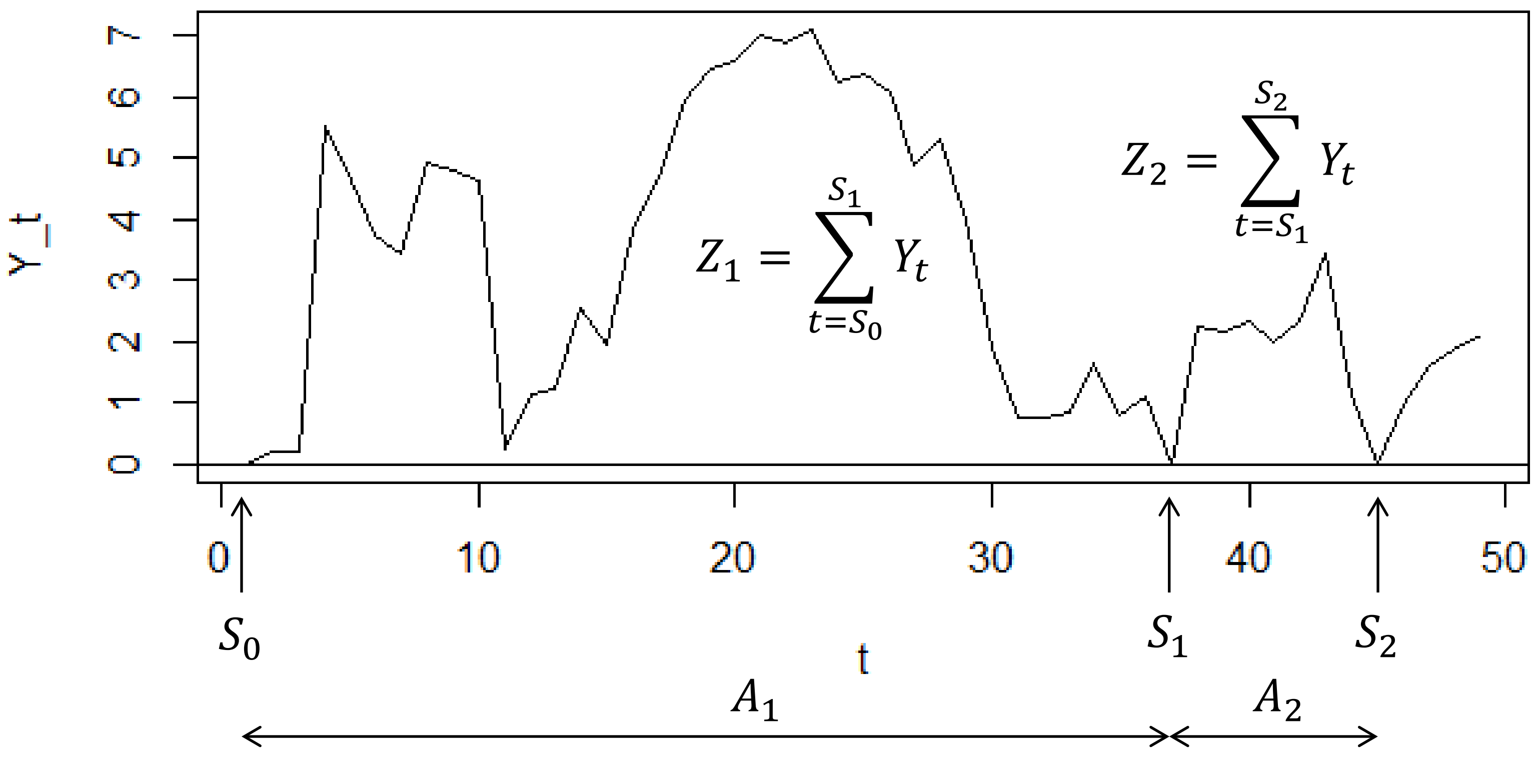

Let the cumulative output during the ith cycle be

$$Z_i = \int_{S_{i-1}}^{S_i} Y_t \, dt$$

The integral becomes a sum for discrete-time output processes.

If Y_t is regenerative, $\mathrm{E}(|Z_1|) < \infty$, $\mathrm{E}(A_1) < \infty$ (and A_1 is aperiodic for discrete-time processes) then

- $Y_t \xrightarrow{D} Y$ as $t \to \infty$
- $\lim_{t\to\infty} t^{-1} \int_0^t Y_s \, ds \xrightarrow{a.s.} \mathrm{E}(Z_1)/\mathrm{E}(A_1)$
- $\mu = \mathrm{E}(Y) = \mathrm{E}(Z_1)/\mathrm{E}(A_1)$

Regenerative structure establishes that there is a steady-state, and the sample mean satisfies a SLLN for the steady-state mean.

Also a CLT?

Let $V_i = Z_i - \mu A_i$. Notice that the V_i are i.i.d.

If it is also the case that $\mathrm{E}(V_1^2) < \infty$, then

- $\lim_{t\to\infty} \sqrt{t}(\bar{Y}_t - \mu) \xrightarrow{D} \gamma\,\mathrm{N}(0,1)$, where $\bar{Y}_t = t^{-1}\int_0^t Y_s\,ds$
- The asymptotic variance $\gamma^2 = \mathrm{E}(V_1^2)/\mathrm{E}(A_1)$.

Therefore, the sample mean satisfies a CLT, and the asymptotic variance γ^2 (more on this later) can be expressed in terms of properties of a regenerative cycle.

Practical value

Establishing the existence of steady state using the regenerative argument means identifying the renewal process and arguing that $\{Y_t; S_i \leq t < S_{i+1}\}$ are i.i.d.

Typically $S_1, S_2, \ldots$ are times such that the state of the simulation is identical, and the probability distributions of all pending events are also the same.

Also need to establish that $\Pr\{A_i < \infty\} = 1$.

Example: For the $M/G/1$ queue let S_i be the ith time an arriving customer finds the system completely empty. Does this work?

Properties of steady-state estimators: Bias

Even though Y_i has a steady state, the estimator $\bar{Y}(m)$ may (usually is) biased. Define the *asymptotic bias*

$$\begin{aligned}\beta &= \lim_{m\to\infty} m\left(\mathrm{E}\left(\bar{Y}(m)\right) - \mu\right) \\ &= \lim_{m\to\infty} m\left(\mathrm{E}\left(\frac{1}{m}\sum_{i=1}^{m} Y_i\right) - \mu\right) \\ &= \sum_{i=1}^{\infty} (\mathrm{E}(Y_i) - \mu)\end{aligned}$$

For the bias in $\bar{Y}(m)$ to disappear, the bias of Y_i must fall off fast enough so that $\beta < \infty$.

Thus, for large m, $\mathrm{Bias}\left(\bar{Y}(m)\right) \approx \beta/m$.

Properties of steady-state estimators: Variance

For any random variables $Y_1, Y_2, \ldots, Y_m$,

$$\mathrm{Var}\left(\bar{Y}(m)\right) = \frac{1}{m^2} \sum_{i=1}^{m} \sum_{j=1}^{m} \mathrm{Cov}(Y_i, Y_j).$$

What we hope in steady-state simulation is that the dependence of $Y_1, Y_2, \ldots$ also stabilizes for large m.

The standard assumption is that beyond some point the process is *covariance stationary*, meaning $\sigma^2 = \mathrm{Var}(Y_m)$ and $\rho_k = \mathrm{Corr}(Y_m, Y_{m+k})$ are no longer a function of m.

$$\mathrm{Var}\left(\bar{Y}(m)\right) = \frac{\sigma^2}{m} \left(1 + 2 \sum_{k=1}^{m-1} \left(1 - \frac{k}{m}\right) \rho_k\right).$$

Asymptotic variance & MSE

For a covariance stationary process the *asymptotic variance* is

$$\gamma^2 = \lim_{m\to\infty} m\mathrm{Var}\left(\bar{Y}(m)\right) = \lim_{m\to\infty} \sigma^2 \left(1 + 2\sum_{k=1}^{m-1}\left(1-\frac{k}{m}\right)\rho_k\right)$$
$$= \sigma^2\left(1 + 2\sum_{k=1}^{\infty}\rho_k\right).$$

For the $\mathrm{Var}\left(\bar{Y}(m)\right)$ to go to 0 we need $\gamma^2 < \infty$.

A combined measure that treats bias and variance equally is the *mean squared error*

$$\mathrm{MSE}\left(\bar{Y}(m)\right) = \mathrm{Bias}^2\left(\bar{Y}(m)\right) + \mathrm{Var}\left(\bar{Y}(m)\right)$$
$$\approx \frac{\beta^2}{m^2} + \frac{\gamma^2}{m}.$$