ZENITH STUDY HUB

Modeling and Simulation

UNIT IV | RANDOM VARIATES

From uniform pseudorandom numbers to credible simulation results

SPPU TE Robotics & Automation | 2024 Pattern | PEC321-ROA (B)

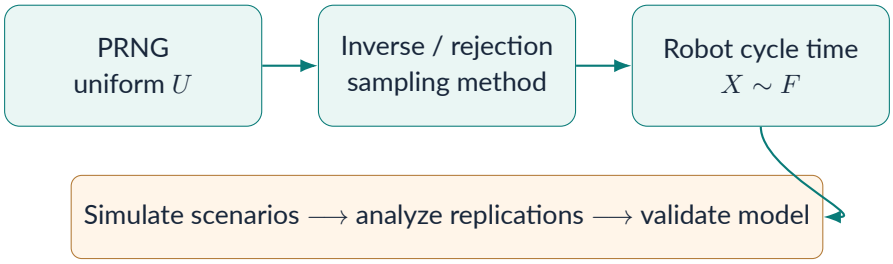

Figure 1. Original end-to-end flow: the random-number generator, sampling algorithm, model inputs, simulation and output-quality checks are separate design decisions.

Primary syllabus
Unit IV, *Random Variate Generation and Output Analysis*, eight hours. Verified from the official SPPU TE Robotics & Automation 2024 Pattern syllabus, course PEC321-ROA (B), physical PDF p. 27. Worked numerical values are constructed for learning, not collected from an industrial robot.

Reviewed edition: 21 November 2026 **Author:** Om Deshmukh

Unit IV | Complete Syllabus Coverage

Official syllabus requirement	Learning material and worked application
Random variate generation	PRNG versus model-specific random variable; quality decisions
Selection of variate generator	Distribution shape, physical support, correlations, reproducibility
Uniform random variables	$X = a + (b - a)U$; numerical and interval examples
Exponential random variables	Inverse CDF with robotic arrivals and units
Normal random variables	Box–Muller transformation and signed-versus-positive warning
Weibull random variables	Shape/scale inverse and lifetime example
Poisson distribution	Discrete event counts, exact inverse-CDF example
Inverse transform	General derivation, steps and boundary conditions
Acceptance–rejection	Majorizing envelope, correct acceptance probability, efficiency
Output data analysis	Mean, percentiles, variability, comparisons and uncertainty
Terminating/non-terminating models	Shift-ending versus long-run factory scenarios
Steady-state analysis	Transient initialization, warm-up and batch means
Confidence interval estimation	Student-t interval with independently replicated means
Verification/validation	Implementation checks versus real-world adequacy
Sensitivity analysis	One-factor and scenario-based robustness calculations

Technical correction to read carefully
The official syllabus lists the Poisson distribution alongside "continuous random variables." Mathematically, **Poisson is a discrete count distribution**. This guide keeps the syllabus topic but teaches the correct classification.

In simple language
A **random number** is usually an approximately uniform U between zero and one. A **random variate** is a simulated value of the uncertain quantity we actually want, such as a repair time, a demand count or an inspection delay. These are different objects.

1 Why Random Variates Are Necessary

1.1 From a clock to a production model

A deterministic spreadsheet may enter *exactly 8 seconds* for every robot operation. Actual operations can vary with product geometry, gripper retries, fixture availability, delays and operator state. A stochastic model replaces some fixed values by carefully justified probability laws. Generating variates from those laws lets us study how output metrics change from run to run.

Input variable	Why its distribution must be justified
Part arrival interval	May be scheduled, clustered or random; Poisson arrivals are not automatic
Robot processing duration	Often positive, product-dependent and not necessarily memoryless
Number of defective parts	Count of events; may show batch dependence or correlated causes
Time to failure	Can depend on wear or prior age; Weibull may be a candidate
AGV travel duration	Can depend on corridor congestion and charging state

Two separate questions
(1) Which distribution represents the measured process? That is *input modeling*, reviewed in Unit III. (2) How does the computer generate values with that distribution? That is *random-variate generation*, the first half of this unit.

1.2 Requirements for selecting a generator

A practical sampling method must give values on the correct support (e.g., non-negative durations), represent the intended probability law, be computationally suitable, allow reproducible experiments and fit the dependencies in the observed system. A fast library function is not a substitute for validated distribution assumptions.

2 Uniform Variates: The Simplest Inverse Transform

2.1 Continuous uniform distribution

For $X \sim \text{Uniform}(a, b)$ with $a < b$,

$$f_X(x) = \frac{1}{b-a} \quad (a \le x \le b), \qquad F_X(x) = \frac{x-a}{b-a}.$$

Set $U = F_X(X)$ and solve for X:

$$\boxed{X = a + (b-a)U}.$$

Numerical 1: conveyor inspection delay
A teaching model uses a delay uniformly distributed from 5 to 13 seconds. For $U = 0.375$,

$$X = 5 + (13-5)(0.375) = \boxed{8.000 \text{ s}}.$$

Check the bounds: a valid draw lies between 5 and 13 seconds. The interpretation is conditional on the uniform model being justified.

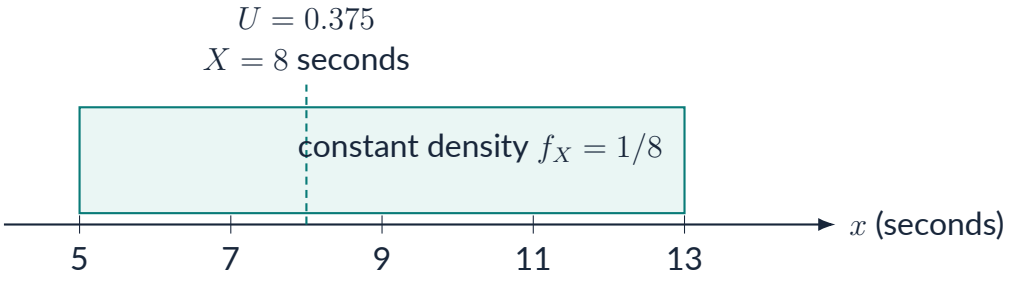

Figure 2. Original uniform-density sketch. The sampled value is determined by its cumulative probability, not by where the diagram is visually tallest.

2.2 Uniform discrete sampling is different

To select one of four equiprobable part types, the numbers can be mapped as $U \in [0, .25), [.25, .50), [.50, .75), [.75, 1)$. Unequal probabilities require unequal cumulative intervals. Floating-point endpoint conventions matter: libraries generally sample in a half-open interval $[0, 1)$, but formula implementations should guard against attempting $\ln(0)$.

3 General Inverse Transform Sampling

Let F_X be a continuous, strictly increasing cumulative distribution function on the intended support. If $U \sim \text{Uniform}(0, 1)$, then

$$\boxed{X = F_X^{-1}(U)}$$

has CDF F_X. For distributions with jumps or a non-strictly increasing CDF, use the generalized inverse $F^{-1}(u) = \inf\{x : F(x) \geq u\}$. This handles discrete counts as well.

1. Derive or look up the CDF $F_X(x)$ and confirm its support and parameter meaning.
2. Generate a suitable independent U value.
3. Solve $F_X(x) = U$ for x, or use a generalized-inverse table.
4. Verify bounds, units and sampling assumptions.
5. Repeat with a suitable random-number stream, preserving its metadata.

3.1 Why the transformation works

For strictly increasing F, $P(X \leq x) = P(F^{-1}(U) \leq x) = P(U \leq F(x)) = F(x)$. The computation preserves the chosen CDF by construction, but does *not* prove the CDF was the correct physical model.

Common numerical error
Inverting F creates a sample. It does not mean $F(x)$ is the probability *density*; it is the probability $P(X \leq x)$. Do not mix a CDF with its derivative, the PDF.

4 Exponential Variates and Interarrival Times

For an exponential distribution with *rate* $\lambda > 0$ (inverse time),

$$F_X(t) = 1 - e^{-\lambda t}, \quad t \geq 0.$$

Set $U = 1 - e^{-\lambda t}$, rearrange $e^{-\lambda t} = 1 - U$, then take logarithms:

$$\boxed{t = -\frac{\ln(1-U)}{\lambda}}.$$

Equivalently, $t = -\ln(U)/\lambda$ for a fresh uniform U: the two formulas have the same distribution, but generally yield different values for the *same* numerical U.

Numerical 2: arrival time
Assume a homogeneous Poisson arrival process with rate $\lambda = 0.25$ per minute (mean interval $1/\lambda = 4$ minutes). If $U = 0.60$,

$$t = -\ln(0.40)/0.25 = \boxed{3.6652 \text{ minutes}}.$$

Check: $1 - e^{-0.25(3.6652)} \approx 0.6000$. The input assumption is explicit and illustrative.

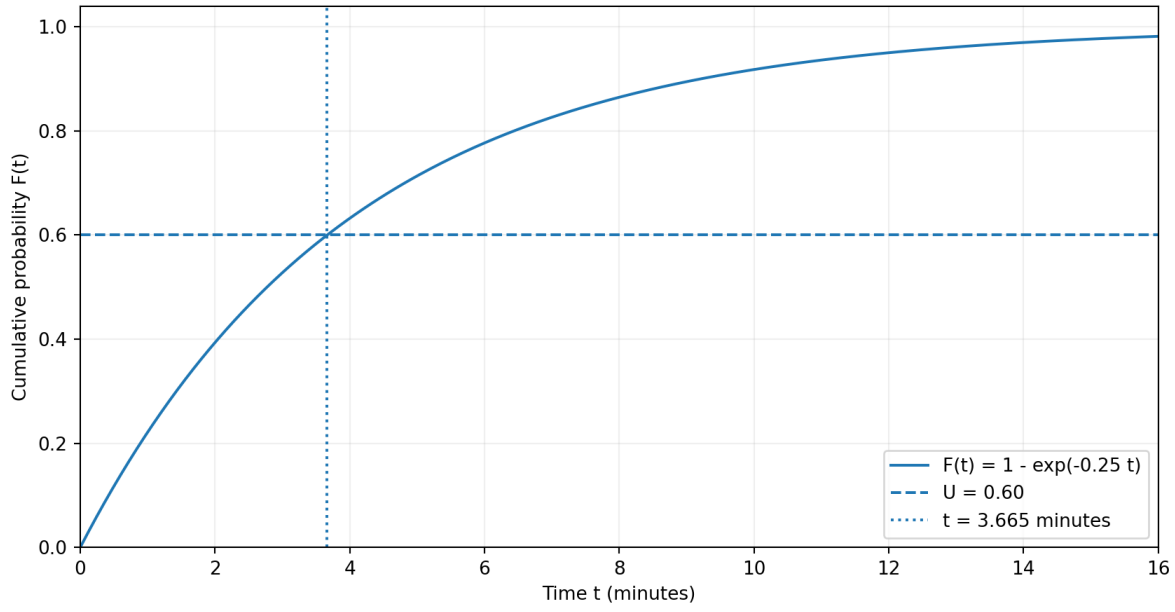

Figure 3. Original exponential CDF and inverse sampling for $U = 0.6$, rate $0.25\,\text{min}^{-1}$. A numerical crosshair identifies the result.

4.1 Mean, variance and applicability

The exponential mean is $E[X] = 1/\lambda$ and variance is $1/\lambda^2$. Its memoryless property states $P(X > s + t \mid X > s) = P(X > t)$. It is useful where its process as-

sumptions are plausible; it is *not* a default model for human-controlled or wear-driven operations. A scheduled inbound truck may violate the homogeneous-arrival assumption.

5 Normal Variates with the Box–Muller Transform

For U_1, U_2 independent and uniform on $(0, 1)$, define

$$Z_1 = \sqrt{-2 \ln U_1}\cos(2\pi U_2), \qquad Z_2 = \sqrt{-2 \ln U_1}\sin(2\pi U_2).$$

Both are independent standard normal variates under ideal input assumptions. A general normal sample follows from $X = \mu + \sigma Z$ for $\sigma > 0$. The formula requires independent draws; do not reuse the same U for U_1 and U_2.

Numerical 3: normal perturbation
Let $U_1 = 0.25$, $U_2 = 0.50$, $\mu = 50$ and $\sigma = 5$. Then $\cos(\pi) = -1$,

$$Z_1 = -\sqrt{-2 \ln(0.25)} \approx -1.6651.$$

Hence $X = 50 + 5(-1.6651) = \boxed{41.6745}$. The outcome unit is whatever μ and σ measure.

Robot-cycle warning
A normal distribution spans negative values. For delays and lifetimes, even a small negative probability can be physically impossible. Evaluate a positive-support model (e.g., lognormal or Weibull) or a properly specified truncation rather than silently replacing negative draws with zero.

6 Weibull Variates for Lifetime Studies

For Weibull *shape* $k > 0$ and *scale* $\eta > 0$,

$$F_T(t) = 1 - \exp\left[-\left(\frac{t}{\eta}\right)^k\right], \quad t \geq 0.$$

Its inverse-transform sampler is

$$\boxed{T = \eta\,[-\ln(1-U)]^{1/k}}.$$

The scale η has the same unit as lifetime; the shape k is dimensionless. For $k = 1$, Weibull reduces to the exponential with rate $1/\eta$. For many applications, $k > 1$ suggests increasing hazard, $k = 1$ constant hazard and $k < 1$ decreasing hazard; the interpretation assumes the Weibull model is appropriate.

Numerical 4: hypothetical belt lifetime
With $\eta = 100$ hours, $k = 2$ and $U = 0.70$,

$$T = 100\sqrt{-\ln(0.30)} = \boxed{109.7257 \text{ hours}}.$$

Interpret this as one random lifetime sample, *not* a prediction that all belts will last approximately 110 hours.

6.1 How to evaluate the model

Inspect maintenance and failure records, censored observations, loading conditions and whether preventive replacement changes the hazard. A fitted Weibull sample cannot substitute for the physical safety and maintenance requirements of a robotic cell.

7 Poisson Random Variates: Discrete Counts

Poisson(θ) models the number N of events in a defined interval with mean $\theta > 0$:

$$P(N = k) = e^{-\theta}\frac{\theta^k}{k!}, \qquad k = 0, 1, 2, \ldots$$

Mean and variance are both θ. If a constant arrival rate is r per hour over duration T hours, the count parameter is $\theta = rT$. Do not mistake r for θ without the observation time.

Numerical 5: inverse CDF for a defect count
Suppose a purely hypothetical batch model uses $N \sim \text{Poisson}(2)$. For $U = 0.70$, cumulative probabilities are

$$F(0) = 0.1353, \quad F(1) = 0.4060, \quad F(2) = 0.6767, \quad F(3) = 0.8571.$$

Find the smallest integer k with $F(k) \geq 0.70$. The answer is $\boxed{N = 3}$ defects.

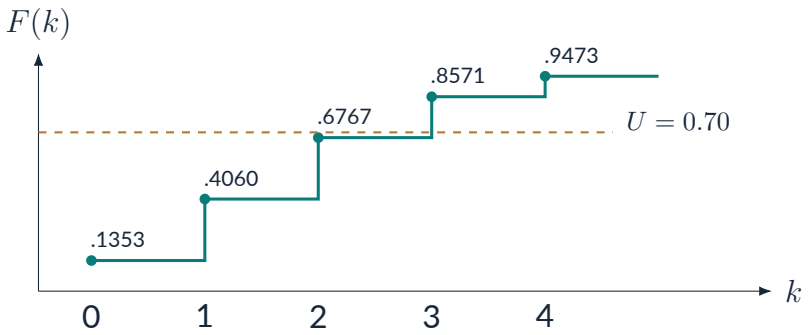

Figure 4. Original step-CDF illustration for Poisson(2). The smallest count with CDF at least 0.70 is 3.

7.1 When Poisson sampling is inappropriate

A Poisson model is not automatically appropriate for clustered defects, product-specific quality problems or an overdispersed count whose variance is much larger than its mean. It is also not a distribution for a positive *duration*; the interarrival times of a homogeneous Poisson process are exponential.

8 Acceptance–Rejection: Sampling Harder Densities

When F^{-1} is unavailable or inconvenient, choose a proposal density $g(x)$ from which sampling is easy. Find a constant $M \geq 1$ satisfying $f(x) \leq Mg(x)$ for all relevant x. Then:

1. Draw a proposed value $Y \sim g$.
2. Draw an independent $V \sim \text{Uniform}(0, 1)$.

3. Accept Y if $V \leq f(Y)/(Mg(Y))$; otherwise discard it and retry.

If the target f and proposal g are normalized probability densities, and the envelope is valid, the unconditional acceptance probability is $1/M$, so expected proposals per accepted sample equal M.

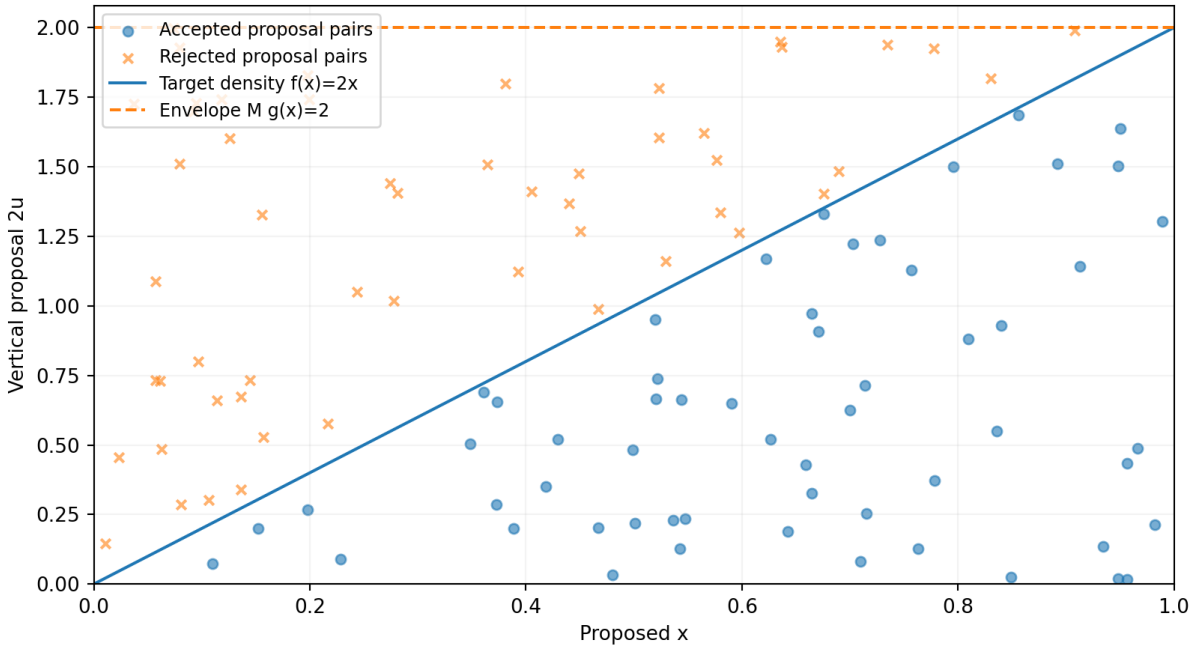

Figure 5. Original rejection-sampling geometry for normalized $f(x) = 2x$ on $[0, 1]$, proposal $g(x) = 1$, and upper envelope $Mg(x) = 2$. Acceptance lies below the target curve.

Numerical 6: accept or reject?
Target $f(x) = 2x$ on $[0, 1]$; proposal $g(x) = 1$; choose $M = 2$. Then $f(x)/(Mg(x)) = x$.
A proposed value $Y = 0.70$ with independent $V = 0.40$ is **accepted** ($0.40 \leq 0.70$).
A proposed value $Y = 0.25$ with $V = 0.80$ is **rejected** ($0.80 > 0.25$).
The exact expected acceptance rate is $1/M = 0.50$; in a small random sample, the observed rate need not equal exactly 50%.

Critical algorithm check
A proposed point must be generated from g, and the acceptance uniform must be independent. Choosing an envelope that fails to bound f produces an invalid sampler. Incorrectly using f/g without the envelope constant can give acceptance probabilities greater than one.

9 Choosing and Implementing a Sampler

Distribution	Simple educational method	Most important check
Uniform(a, b)	Affine transform	$a < b$ and output bounds
Exponential(λ)	Inverse CDF	Rate versus mean (scale) units
Normal(μ, σ)	Box–Muller	$\sigma > 0$ and physical support
Weibull(k, η)	Inverse CDF	Shape/scale roles and unit of η
Poisson(θ)	Discrete inverse CDF	Count-valued output; defined interval
Complex density	Rejection sampling	Valid envelope Mg and independent draws

In production, prefer well-tested library functions and documented pseudorandom streams over reinventing distribution samplers. A classroom algorithm can be mathematically correct yet much slower or numerically weaker than a mature implementation.

9.1 Reproducible NumPy demonstration

Listing 1: Original example: draw selected model variates reproducibly

```
import numpy as np
rng = np.random.default_rng(20261121)
arrivals = rng.exponential(scale=4.0, size=5)  # minutes
lifetimes = 100.0 * rng.weibull(a=2.0, size=5)  # hours
counts = rng.poisson(lam=2.0, size=5)
print(arrivals, lifetimes, counts)
```

NumPy parameter caution

The NumPy exponential API uses a *scale* equal to the mean $1/\lambda$. Its Weibull sampler returns the dimensionless standard-scale form; multiply by η to obtain the physical lifetime scale. Record NumPy version and bit-generator details to reproduce an experiment accurately.

10 From Simulation Runs to Output Analysis

Outputs are random measurements produced by repeated *simulation experiments*, for example jobs completed per shift, fraction of time blocked, mean job waiting time, 95th-percentile queue delay and number of rework attempts.

One simulation run is one realization. An appealing single animation or a single run average cannot establish the typical output with high certainty. A final numerical result needs a defined horizon and a method for quantifying variability.

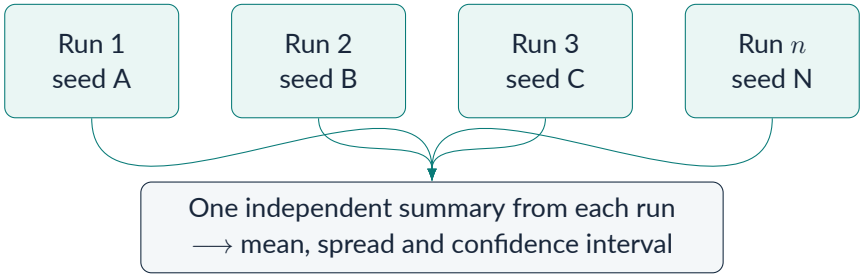

Figure 6. Original independent-replication scheme. Each output summary represents one replication under the same model conditions, not a new physical observation from a factory.

10.1 A checklist for reporting an output

State the metric and units, experiment starting condition, warm-up choice if any, simulation end condition, independent replications and stream management, sample size, uncertainty interval and noteworthy high-impact scenarios.

10.2 Descriptive summaries

For replication outputs $Y_1, \ldots, Y_n$, estimate $\bar{Y} = \sum_i Y_i/n$. A median describes the central ranking, while a 90th/95th percentile emphasizes tail delays. They cannot be interchanged. For utilization, give the time interval and the denominator: busy time divided by available time is not identical to productive time divided by shift length if maintenance is excluded.

11 Terminating versus Non-terminating Simulations

11.1 Terminating model

A terminating experiment has a meaningful finite end: one shift, one batch, a specific repair operation or a fixed customer order. The system's startup conditions may be physically important, so deleting the opening period arbitrarily can remove real effects.

Robot example: Compare two grippers over 8-hour shifts, starting each morning with an empty buffer. Record completed parts, rejects and overtime with consistent initial conditions.

11.2 Non-terminating / steady-state model

A non-terminating model represents ongoing operation where long-run behavior is of interest. Start-up transients can bias a long-run mean. A warm-up period may be discarded when estimating steady-state behavior, but the warm-up choice requires evidence and a relevant convergence diagnostic.

Robot example: Evaluate long-run mean waiting time for an automated loading station across many operating days, under stable demand assumptions. If demand changes by shift or season, there may be no single stationary steady state to estimate.

Feature	Terminating	Steady-state
Natural horizon	Finite shift/batch	Long-run regime
Startup data	Part of the target process	May contain transient bias
Common analysis	Independent runs	Delete warm-up; batch means
Typical objective	Shift throughput	Long-run mean delay

12 Warm-up, Replication-Deletion and Batch Means

12.1 Initial transient

If a queue begins empty, the first jobs may have very short waits. Those early low waits could bias an estimator for the long-run, routinely congested cell. Conversely, an initially full buffer could produce a transient high-wait bias.

In replication-deletion, run independent long replications and discard a justified initial warm-up portion t_w from each. Measure the average over $[t_w, T]$. Never choose t_w only to obtain a desired conclusion.

12.2 Batch means for a long run

Partition a sufficiently long steady-state segment into non-overlapping batches and compute one mean per batch. Large batches can reduce dependence between batch means; small batches may retain substantial serial correlation. Their approximate independence and adequate number should be diagnosed before treating them like independent observations.

Statistical warning
Consecutive customers *inside one queue run* are often correlated. A naive t interval built by treating every customer waiting time as an independent observation can seriously understate uncertainty. Either use independent replication-level summaries or an appropriate time-series/batch-means method.

13 Confidence Interval for the Mean Output

For independent replication summaries Y_i under identical model conditions, the usual two-sided $(1-\alpha)100\%$ Student-t interval is

$$\boxed{\bar{Y} \pm t_{1-\alpha/2,n-1}\frac{s_Y}{\sqrt{n}}}, \qquad s_Y = \sqrt{\frac{\sum_{i=1}^{n}(Y_i-\bar{Y})^2}{n-1}}.$$

An exact small-sample t interval needs normality of the replication summaries and independence; approximation may be useful for sufficiently large sample sizes under suitable conditions. A confidence interval quantifies uncertainty in an estimated *mean*, not the range containing 95% of individual robot delays.

Numerical 7: five independent simulation replications

Suppose five hypothetical shifts yield output averages $28, 30, 32, 29, 31$ (minutes of mean delay per replicated experiment). Then

$$\bar{Y} = 30, \qquad \sum_i (Y_i - \bar{Y})^2 = 4 + 0 + 4 + 1 + 1 = 10.$$

$$s_Y^2 = 10/(5-1) = 2.5.$$

$$s_Y = \sqrt{2.5} = 1.58114, \qquad t_{0.975,4} = 2.77645.$$

$$h = 2.77645 \frac{1.58114}{\sqrt{5}} = 1.96324.$$

Thus the illustrative 95% interval is $\boxed{(28.04, 31.96)}$ minutes, conditional on assumptions above.

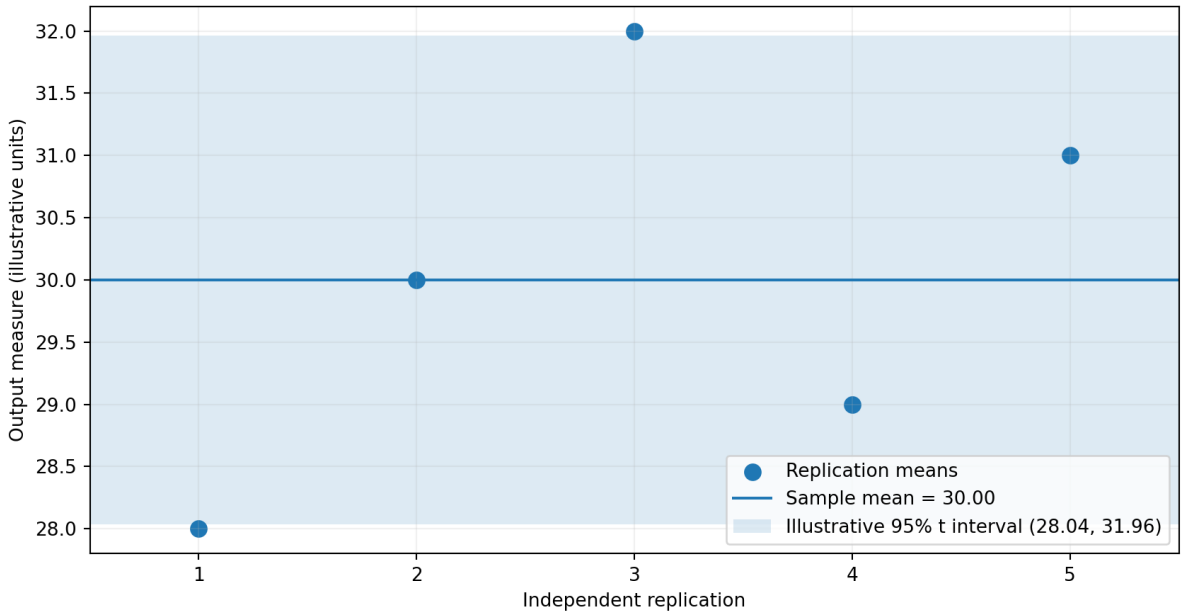

Figure 7. Original five-replication summary and Student-t confidence interval. The center and half-width are calculated from the five shown values, not from a real cell.

13.1 Interpretation and precision

If the confidence procedure were repeated across many independently generated datasets from an appropriate model, approximately 95% of those constructed intervals would contain the true mean. The confidence interval does

not assign a 95% probability to a single fixed unknown parameter in the usual frequentist interpretation.

For comparable variance and assumptions, quadrupling the number of independent replications reduces the standard error by about a factor of two. It does not fix biased inputs, an incorrect model or correlated replications.

13.2 Comparing competing cell designs: paired replications

Comparing two designs with the *same input stream* can reduce random variability when a fair correspondence exists. Use one independent seed for each *pair* of runs; apply the same arrivals and service-time draws to designs A and B, then compare the *paired difference* $D_i = Y_{B,i} - Y_{A,i}$. Pairs must be independent of one another, even though the two outputs *within* a pair may be correlated.

Numerical 9: hypothetical change in completed parts per shift
Suppose five paired differences are $D = [5, 7, 4, 6, 8]$ parts. Then $\bar{D} = 6$, $s_D^2 = 2.5$, and for $n = 5$ independent pairs the 95% Student-t interval is

$$\bar{D} \pm t_{0.975,4}\frac{s_D}{\sqrt{5}} = 6 \pm 1.9632 = \boxed{(4.04, 7.96) \text{ parts}}.$$

Under the stated sampling assumptions this estimates a positive average difference; it is *not* proof the real factory will produce exactly six extra parts.

Comparison trap
Never subtract the endpoints of two separately reported confidence intervals and call the result the confidence interval for a design difference. Analyze replicated differences directly if the runs are paired, or use an appropriate unpaired comparison method otherwise.

14 Sensitivity Analysis and Engineering Decisions

Sensitivity analysis changes an assumed input or a controllable parameter while keeping the rest of the experiment interpretable. Analyze whether a decision

changes when the cycle time, demand, downtime or failure model varies plausibly.

Numerical 8: ideal throughput sensitivity
In an ideal, single-server processing station with no downtime, blocking or changeovers, maximum theoretical output per hour is $3600/c$ if each cycle takes c seconds:

c (s/part)	6	8	10
Ideal rate (parts/h)	600	450	360

The 8-to-10-second change reduces the ideal bound by $450 - 360 = \boxed{90}$ parts/hour or 20% of its initial value. Actual factory throughput can be lower.

14.1 One-factor versus scenario analysis

One-factor: vary one cycle-time parameter while other assumptions are fixed. **Scenario:** combine changes such as high demand, more failures and slower repairs to represent a plausible stressful operating state. Correlated inputs should not be artificially varied independently if real physics couples them.

Practical decision rule
A recommendation is more credible when its ranking survives reasonable changes in input assumptions and replication uncertainty. If a tiny change reverses the preferred design, report the decision as fragile, not certain.

15 Verification, Validation and Reproducibility

Verification tests whether program logic and formulas implement the specified conceptual model. **Validation** tests whether the model is a credible representation of the real system for a stated decision question. A visually realistic robot animation can pass neither.

1. **Unit-test the samplers:** inverse CDFs produce correct support and sample

means; rejection envelopes bound target densities.

2. **Test event logic:** run deterministic arrivals and service times against a hand-computed small example.
3. **Inspect real measurements:** compare cycle-time, downtime and arrival distributions with observed operations.
4. **Check output measures:** compare model throughput/queue metrics with a held-out observed shift or other appropriate evidence.
5. **Record assumptions:** distribution families, parameters, initial state, warm-up policy, scenario definitions and confidence method.
6. **Preserve experiments:** seeds, bit-generator identifiers, software versions, code revision and unit conventions.

Evidence boundary
Passing arithmetic tests cannot certify a random-variate model for a specific manufacturing plant. No plant observations, safety certification, or physical hardware runs have been supplied for these worked examples.

16 Exam-Ready Questions and Worked Answers

1. **Define random variate.** A sample of an uncertain model variable with a specified distribution, e.g., a repair duration.
2. **How does random number differ?** A typical PRNG produces uniform-like numbers; a sampling method maps them into model-specific variates.
3. **Give the inverse transform.** $X = F^{-1}(U)$, using generalized inverse for discrete CDFs.
4. **Uniform inverse numerical:** $[5, 13]$, $U = .375$ yields 8 seconds.
5. **Exponential numerical:** rate .25/min and $U = .6$ yields 3.6652 minutes.
6. **Normal sampling:** use two independent uniforms in Box–Muller; $X = \mu + \sigma Z$.
7. **Weibull parameters:** k shape dimensionless, η scale with time unit.

8. **Weibull numerical:** $k = 2, \eta = 100, U = .7$ gives 109.7257 hours.
9. **Is Poisson continuous?** No: it is an integer-valued count distribution; the syllabus listing is misleading.
10. **Poisson inverse numerical:** $\theta = 2, U = .7$ selects $k = 3$.
11. **What does acceptance–rejection require?** An easy proposal g, a valid constant M with $f \leq Mg$ and an independent accept/reject draw.
12. **Expected acceptance rate?** $1/M$ for normalized densities.
13. **Explain terminating simulation.** It has a natural experiment end such as one production shift.
14. **Why use warm-up?** To reduce start-up transient bias when estimating a long-run steady-state mean.
15. **What is a replication?** An independently randomized complete simulation run with the same intended model scenario.
16. **Confidence interval example:** five replicate means $[28, 30, 32, 29, 31]$ produce a 95% Student-t interval $(28.04, 31.96)$ under stated conditions.
17. **What is sensitivity analysis?** Exploring how scenario results change with plausible parameters and assumptions.
18. **Verification versus validation?** Correct implementation versus adequate real-system representation for an intended use.
19. **How can samples be reproducible?** Save algorithm/bit-generator, software version, seed, code and sampling parameters.
20. **Why are confidence intervals not a cure for bad models?** They quantify sampling uncertainty conditional on the model, not model discrepancy or invalid data.

Sources, Provenance and Review Limits

Required syllabus: SPPU TE Robotics & Automation (2024 Pattern), PEC321-ROA (B), Unit IV, *Random Variate Generation and Output Analysis*, physical PDF page 27.

Available campus material: The authenticated *ENDSEM NOV 2026* Drive collection contains M&S Unit I/II slides and question banks, the syllabus, and a

Jerry Banks textbook. Those Unit I/II files were inspected for course framing but are *not* claimed to contain a Unit IV lesson. The Unit IV lecturer handout was not available in the folders checked.

Extracted scholarly reference: Manuel D. Rossetti, *Simulation Modeling and Arena*, open 3rd edition; reviewed the appendix on pseudo-random numbers and random variates (physical PDF pp. 764–801), and the chapter on statistical analysis of infinite-horizon models (pp. 397–455) in the accessible local reference copy. Textbook images have not been copied.

External references: NumPy Random Generator API; SAS Simulation Studio output analysis; NIST guidance on confidence limits for means; Winter Simulation Conference: batch means and independent replications.

Images: Three original mathematical plots prepared for this unit, plus original TikZ vector diagrams; no third-party photo or textbook diagram has been embedded. Source PNG/SVG files, formulas and verification code are included in the downloadable QA package.

Limitations: Mathematical examples use invented learning data. Source citations do not imply measurement-based calibration or certified manufacturing performance. This v2 review adds an explicitly worked paired-scenario comparison and its uncertainty limits. The PDF remains a review candidate pending independent academic sign-off and real simulation validation.

UNIT IV | P0 syllabus-gap corrections

Supplement to the existing reviewed edition. All original teaching pages remain unchanged. The material below closes only the specifically rejected syllabus/question-bank gaps.

1. Departmental Q4: Five exponential random variates

Given exponential mean = 5 time units, rate λ = 1/5, and uniform draws U = (0.10, 0.45, 0.62, 0.81, 0.93). From F(x)=1−exp(−x/5), set U=F(x) and solve X=−5 ln(1−U).

U	1−U	Expression	X (time units)
0.10	0.90	−5 ln(0.90)	0.5268
0.45	0.55	−5 ln(0.55)	2.9892
0.62	0.38	−5 ln(0.38)	4.8379
0.81	0.19	−5 ln(0.19)	8.3037
0.93	0.07	−5 ln(0.07)	13.2963

These are individual sampled variates, not the sample mean. The alternative −5 ln(U) produces the same distribution when U is freshly uniform, but not the same numerical result for these fixed random digits.

2. Departmental Q12: Exponential variate with mean 4

With 1/λ=4, λ=0.25 per time unit and U=0.7: X=−ln(1−0.7)/0.25 = −4 ln(0.3) = 4.8159 time units.

3. Departmental Q13–Q14: Weibull variates

For scale α>0, shape β>0: F(x)=1−exp[−(x/α)^β] for x≥0. Rearranging gives X=α[−ln(1−U)]^(1/β). The scale carries the time unit; the shape is dimensionless.

Question	Given	Substitution	Result
Q13	α=3, β=2, U=0.6	3[−ln(0.4)]^(1/2)	2.8717
Q14	α=4, β=1.5, U=0.75	4[−ln(0.25)]^(2/3)	4.9731

4. Departmental Q11: Triangular inverse-transform derivation

For minimum a, mode b and maximum c with a<b<c, the triangular density is f(x)=2(x−a)/[(b−a)(c−a)] for a≤x≤b, and f(x)=2(c−x)/[(c−a)(c−b)] for b<x≤c. Its CDF is F(x)=(x−a)²/[(b−a)(c−a)] for a≤x≤b and F(x)=1−(c−x)²/[(c−a)(c−b)] for b<x≤c.

At the mode, F(b)=(b−a)/(c−a)=p. For U≤p solve U=(x−a)²/[(b−a)(c−a)] to get X=a+sqrt[U(b−a)(c−a)]. For U>p solve 1−U=(c−x)²/[(c−a)(c−b)] to get X=c−sqrt[(1−U)(c−a)(c−b)]. Thus both branches join at x=b without a jump.

Check using a=1, b=3, c=7, so p=2/6=1/3: U=1/6 yields X=1+sqrt(2)=2.4142; U=2/3 yields X=7−sqrt(8)=4.1716. Both lie on their respective sides of the mode. For a boundary-mode triangular distribution (a=b or b=c), use the continuous limiting expression rather than dividing by zero.

Question-bank alignment: M&S Question Bank.pdf, Unit IV, Q4 and Q11–Q14. These worked cases supplement, rather than replace, the source PDF’s existing exponential and Weibull examples.