

Bayesian vs. Frequentist SPC: When the Prior Matters

Posterior distributions for process parameters. Credible intervals vs. confidence intervals. When informative priors from process history change the capability decision, and when 200 observations make the distinction irrelevant.

Ref: Box & Tiao · Hamada & Johnson · ISO 22514-2 · NIST/SEMATECH Handbook

I — Publication

Bayesian vs. Frequentist SPC: When the Prior Matters

Posterior distributions for process parameters. Credible intervals vs. confidence intervals. When informative priors from process history change the capability decision, and when 200 observations make the distinction irrelevant.

Ref: Box & Tiao · Hamada & Johnson · ISO 22514-2 · NIST/SEMATECH Handbook

Date 2026-02-10

Author Eric Wolters

Reading 18 min

Bayesian · SPC

II — The Frequentist Framework

Shewhart's Foundation

Walter Shewhart's framework, published in 1931, treats process parameters as fixed but unknown quantities. The sample mean and sample standard deviation are point estimates. Control limits are set at ± 3 standard deviations from the process center, derived from a historical baseline period assumed to represent the in-control state. A point beyond the limits triggers investigation. The framework makes no probability statement about the process parameters themselves—only about the likelihood of observing the data under the null hypothesis that the process has not changed.

What "In Control" Means Statistically

A process is in statistical control when all observed variation is attributable to a constant system of chance causes. Operationally, this means the control chart shows no points beyond the limits and no non-random patterns. The frequentist interpretation is a hypothesis test: we fail to reject the null hypothesis that the process is operating in its baseline state. This does not mean the process is good,

centered, or capable. It means only that it is stable—predictable. Capability analysis begins where control chart analysis ends.

Confidence Intervals for Cpk

A frequentist 95% confidence interval for Cpk does not mean "there is a 95% probability that the true Cpk lies in this interval." It means: if we repeated this sampling procedure many times, 95% of the constructed intervals would contain the true value. The distinction is subtle but consequential. No single interval has a 95% probability of containing the truth. The probability statement is about the procedure, not about the parameter. For a plant manager asking "what is the probability our Cpk exceeds 1.33?", the frequentist framework does not directly answer the question.

III — The Bayesian Framework

Posterior Distributions for Process Parameters

The Bayesian approach treats process parameters (μ, σ) as random variables with probability distributions that encode our state of knowledge. Before collecting data, we specify a prior distribution reflecting what we know (or don't know) about the process. After observing data, Bayes' theorem updates the prior to a posterior distribution. Every quantity derived from the parameters—including Cpk—inherits a posterior distribution.

The posterior distribution of Cpk directly answers probabilistic questions: $P(\text{Cpk} > 1.33 \mid \text{data}) = 0.82$ means, given the observed data and the prior, there is an 82% probability that the true Cpk exceeds 1.33. This is the statement the plant manager actually wants. No reinterpretation required.

Credible Intervals vs. Confidence Intervals

A 95% credible interval is an interval that contains the true parameter with 95% posterior probability. This is the intuitive interpretation that most practitioners incorrectly assign to confidence intervals. The Bayesian credible interval delivers what the confidence interval is commonly (but wrongly) believed to deliver.

For large samples, credible intervals and confidence intervals are numerically similar. The practical difference emerges with small samples, where the prior has influence, and in interpretation, where the credible interval supports direct probability statements about the parameter.

The Computational Mechanism

For normal data with unknown mean and variance, the standard Bayesian model uses a normal-inverse-chi-squared conjugate prior. The posterior for (μ, σ^2) is available in closed form, and the posterior for Cpk is obtained by transformation. In practice, this means drawing samples from the posterior of (μ, σ) and computing Cpk for each draw, producing a Monte Carlo approximation to the posterior of Cpk. The [Bayesian Cpk calculator](#) performs this computation automatically.

IV — Prior Selection

Non-Informative (Reference) Priors

The Jeffreys prior for the normal model is $p(\mu, \sigma) \propto 1/\sigma$. It encodes maximal ignorance about the parameters and lets the data dominate the posterior. For most capability studies, this is the appropriate default. The resulting posterior depends only on the sufficient statistics (sample mean, sample variance, sample size) and introduces no external information. The Bayesian analysis becomes a reparametrization of the frequentist analysis, with the advantage of direct probability statements.

Weakly Informative Conjugate Priors

A weakly informative prior encodes soft constraints: the process mean is probably within the specification range, the standard deviation is probably less than half the tolerance. These constraints prevent the posterior from placing substantial probability on physically implausible parameter values without meaningfully influencing the result when data are abundant. For $n > 30$, a weakly informative prior and the Jeffreys prior produce posteriors that are visually indistinguishable.

Informative Priors from Process History

When substantial historical data exist from the same process under the same conditions, an informative prior can incorporate that knowledge. This is where the Bayesian framework offers a genuine advantage beyond interpretation. Consider a process with 500 historical observations establishing $Cpk \approx 1.5$. A new production run produces 20 observations with sample $Cpk = 1.1$. The frequentist analysis sees only the 20 observations and reports $Cpk = 1.1$ with a wide confidence interval. The Bayesian analysis with an informative prior from the historical data produces a posterior that pulls toward the historical value, reporting a posterior mean near 1.35 with a narrower credible interval.

This is appropriate when the process has not fundamentally changed. It is dangerous when the process has changed and the historical data no longer represent the current state. Prior sensitivity analysis—repeating the analysis with different priors and comparing the posteriors—reveals how much the conclusion depends on the prior versus the current data.

Prior Sensitivity Analysis

Run the analysis with three priors: (1) non-informative, (2) weakly informative, (3) informative from history. If all three posteriors agree on the capability decision, the prior does not matter for this dataset and sample size. If they disagree, the sample is too small relative to the decision boundary, and the prior is doing real work. Report this explicitly. Do not hide the prior dependence behind a single number.

V — Worked Examples

Small Sample (n = 25): Large Divergence

A machining process produces shaft diameters with LSL = 24.95 mm and USL = 25.05 mm. Twenty-five parts are measured: sample mean = 25.003 mm, sample standard deviation = 0.0082 mm. The traditional Cpk is $\min((25.05 - 25.003) / (3 \times 0.0082), (25.003 - 24.95) / (3 \times 0.0082)) = \min(1.91, 2.15) = 1.91$. This looks excellent.

The Bayesian posterior (Jeffreys prior) has mean Cpk = 1.78, median = 1.75, and 90% credible interval [1.24, 2.44]. The probability $P(\text{Cpk} > 1.33) = 0.91$. The point estimate of 1.91 overstates the certainty. There is a 9% chance the process does not meet Cpk = 1.33—not negligible for a safety-critical application.

Large Sample (n = 200): Convergence

Same process, same parameters, but now 200 observations. Sample mean = 25.002 mm, sample standard deviation = 0.0085 mm. Traditional Cpk = 1.88. Bayesian posterior mean = 1.86, 90% credible interval [1.70, 2.03]. $P(\text{Cpk} > 1.33) = 0.999$. The two methods agree. The sample is large enough that the prior is irrelevant and the estimation uncertainty is small. The Bayesian analysis adds nothing actionable.

Borderline Cpk: When the Method Changes the Decision

Injection molding process, 35 observations. Sample mean = 12.48, USL = 12.60, LSL = 12.20. Sample standard deviation = 0.030. Traditional Cpk = $\min((12.60 - 12.48) / 0.090, (12.48 - 12.20) / 0.090) = \min(1.33, 3.11) = 1.33$. The process exactly meets the 1.33 requirement by the point estimate.

Bayesian posterior (Jeffreys prior): mean Cpk = 1.27, 90% credible interval [0.99, 1.59]. $P(\text{Cpk} > 1.33) = 0.43$. The process has less than a coin-flip chance of actually meeting the requirement. The frequentist point estimate says "pass." The Bayesian posterior says "insufficient evidence." For this sample size and this result, the method changes the decision. This is the scenario where the Bayesian approach earns its computational cost.

VI — Practical Guidance

Decision Tree

Use the Bayesian approach when: (1) $n < 50$ and the capability decision matters, (2) the point estimate is within 0.3 of the acceptance threshold (borderline), or (3) the application is high-consequence (aerospace, medical device, automotive safety) and you need a probability statement for the capability claim. Use the traditional approach when: $n > 100$ and the point estimate is far from the threshold, or when the customer specification explicitly requires traditional indices per AIAG or PPAP requirements.

Implementation Costs

The computational cost is negligible. The real cost is organizational: introducing Bayesian methods requires training the quality team on interpretation, updating standard operating procedures, and potentially negotiating with customers who expect traditional Cpk values. Start by running both analyses side by side. Report the traditional index for contractual compliance and the Bayesian posterior for internal decision-making. When the two agree, the traditional report is sufficient. When they disagree, the Bayesian posterior reveals the uncertainty that the point estimate hides.

Explaining Posteriors to a Plant Manager

Do not say "posterior distribution." Say: "Based on 35 measurements, there is a 43% chance our process actually meets the 1.33 Cpk requirement. The traditional calculation says 1.33 exactly, but that number has a margin of error. We need 80 more parts to be confident." The probability statement is what makes the Bayesian result communicable. It converts a statistical abstraction into a risk statement that drives resource allocation.

The [Bayesian Cpk calculator](#) reports both the traditional index and the full posterior summary. The [standard Cpk calculator](#) includes confidence intervals as a frequentist alternative. Use whichever matches your audience and contractual requirements.

VII — Related Instruments

INSTRUMENT	SPECIFICATION
Bayesian Cpk Calculator	P(Cpk > 1.33) and credible intervals. Posterior distribution of capability.
Cpk / Ppk Calculator	Traditional process capability indices. Cp, Cpk, Pp, Ppk, sigma level, DPMO.
Control Chart Generator	I-MR, X-bar/R, X-bar/S, p, c, u charts with Nelson rules.

[← All Whitepapers](#)

SVEND is the consulting practice of Eric Wolters – operational excellence, quality management and strategic planning. ILSSI accredited, SCMEP Technical Partnership Provider.
svend.ai/whitepapers/bayesian-vs-frequentist-spc/ · published 10 February 2026