

Parametric Should-Cost Model V2

Defensible price estimation for machined aerospace parts

Kenyon Woodley | Western Washington University | 2026 | Stack: Python + cadquery + Streamlit + Excel

01 – PROBLEM

Without a cost model, every supplier quote is a guess you can't challenge.

In aerospace manufacturing, sourcing engineers receive supplier quotes and are expected to evaluate them. But evaluation without an independent cost estimate is not evaluation; it is acceptance. The sourcing professional has no leverage, no floor to negotiate from, and no way to explain to leadership why a quote is too high or surprisingly low.

Commercially available should-cost tools exist (aPriori, Costimator, SEER-MFG) but they require engineering data, trained analysts, and enterprise licensing. They are not accessible to a sourcing team making fast portfolio decisions across dozens of parts.

"What's your should-cost on this part?" A question that should take minutes to answer, not days. Rarely had a structured answer at all.

02 – INSIGHT

You don't need geometry to estimate cost. You need the right parameters.

Feature-based costing tools like aPriori read CAD files and count holes, pockets, and tolerances. That precision is valuable, but it requires tooling that most sourcing teams don't have at the point of negotiation. The vast majority of sourcing decisions are made from a drawing title, a material, and a complexity characterization, not a fully instrumented CAM file.

A parametric model trades geometric precision for speed and accessibility. Given material, part weight, complexity tier, tolerance tier, volume, and region, the model produces a defensible price band: not a single number, but a structured range that tells the sourcing engineer where the floor is, where the ceiling is, and which cost driver dominates the estimate.

The model doesn't replace aPriori. It fills the gap between "I have no idea" and "I have a routing." That gap is where most sourcing conversations happen.

03 – APPROACH

A modular Python engine with a traceable, deterministic cost model

The engine is built in Python with a strict separation of concerns: the cost engine produces all numbers deterministically, a separate AI layer generates plain-language narrative from those numbers, and an Excel exporter surfaces everything in a format a sourcing team can actually use in a meeting. Every number traces back to an assumption. No black boxes.

BUILD ORDER — COST ENGINE

$$\text{Unit Price} = (\text{Material} + \text{Machining} + \text{Setup} + \text{Outside Processes} + \text{Scrap}) \times (1 + \text{Overhead Rate}) \times (1 + \text{Margin Rate})$$

Three scenarios computed in parallel: Low (optimistic rates + margin), Mid (baseline), High (conservative)

Overhead basis configurable: all direct cost (default) or labor + machine only

Price floor enforced by complexity tier: low band cannot fall below a defensible minimum

Sensitivity: top 3 cost drivers identified by +10% perturbation analysis

Confidence: demerit-based scoring (High / Medium / Low) driven by input quality and material volatility

The engine is fully modular. Data files (material prices, machine rates, buy-to-fly ratios, outside process costs) live in CSV and JSON, editable without touching code. All assumptions are surfaced in the Excel output so every number is auditable.

V2 added STEP geometry ingestion via cadquery, auto-populating bounding box, volume, and buy-to-fly ratio from uploaded part files, and a Streamlit UI wrapper that puts the full engine in a browser without requiring Python.

04 — MODEL COMPONENTS

Eight cost elements, each traceable to a configurable assumption

01

MATERIAL

Blank weight = finished weight × buy-to-fly ratio. BTF sourced from an 8-material × 5-tier table. Material price per lb in low/mid/high bands. Dominant cost element for heavy, high-BTF parts.

02

MACHINING

Machine rate × machining hours. Hours estimated from complexity tier (1–5) and tolerance tier (standard / tight). Machine type (3-axis, 4-axis, 5-axis, turning) assigned by complexity tier. Dominant driver for most machined parts.

03

SETUP

Setup hours × machine rate, amortized over batch size. Larger batches reduce setup cost per piece. Complexity tier 5 carries 4× the setup hours of tier 1.

04

OUTSIDE PROCESSES

Optional stack: heat treat, anodize, hard anodize, chromate, nickel plate, passivation, shot peen, CMM inspection, paint. Each carries independent low/mid/high cost bands. Per-piece costs. Batch economics are a noted V2 item.

05

SCRAP

Applied to direct cost as a scrap factor (low: 1%, mid: 3%, high: 7%). Surfaced as an explicit line item, not hidden inside overhead. Varies across the price band.

06

OVERHEAD

Applied as a rate on top of direct cost. Basis is configurable: all direct cost (default, conservative, treating overhead as full SG&A proxy) or labor + machine only (treating material as a passthrough). US default: 35%.

07

MARGIN

Supplier margin applied to total cost. Low/mid/high bands by region. US: 8% / 14% / 22%. Asia carries a wider margin band (10% / 18% / 28%) reflecting higher variance in supplier tier and commercial terms.

08

PRICE BAND

Three scenarios computed in parallel using low, mid, and high assumptions across material price, machine rate, scrap factor, and margin. Floor check enforced. Sensitivity identifies top 3 cost drivers by dollar impact.

A range with a rationale, not a number with false precision

The model produces three scenarios simultaneously: Low (optimistic assumptions: best-case material pricing, lowest machine rates, tightest margin), Mid (baseline assumptions, the primary reference point), and High (conservative assumptions: worst-case on all four variable inputs). The band is not a confidence interval in the statistical sense. It is a structured range that reflects real variance in the four inputs that move most: material price, machine rate, scrap rate, and supplier margin.

In leadership reviews, the instinct is to anchor on the mid number. The right framing is different: the low is the floor to negotiate toward, the high is the ceiling that justifies walking away, and the mid is the starting position for a supplier conversation. **The band is the tool. The mid is a reference point, not a target.**

WHAT MOVES ACROSS THE BAND

Low: price_low/lb + rate_low/hr + scrap 1% + margin 8% | Mid: price_mid/lb + rate_mid/hr + scrap 3% + margin 14% | High: price_high/lb + rate_high/hr + scrap 7% + margin 22%

Machining hours and setup hours are held constant across all three scenarios; they are geometry-driven, not market-driven.

Outside process costs have independent low/mid/high values and vary across the band.

Overhead rate is region-configurable (US: 35%, EU: 38%, Asia: 22%) but held constant within a single estimate's three scenarios, reflecting facility structure rather than market volatility.

Band width narrows as outside process and material cost become a larger share of total cost.

Which lever matters most, before the negotiation starts

For every estimate, the engine perturbs four key variables by +10% and measures the resulting change in the mid unit price. The top three drivers are returned in rank order. This answers the most important pre-negotiation question: where is the cost actually coming from?

For machined parts, machine rate and machining hours dominate almost universally. Material price becomes a significant driver only when buy-to-fly is high or material is inherently expensive (Ti, Inconel). Knowing this in advance changes the negotiation strategy.

VARIABLE	WHAT IT REPRESENTS	SOURCING IMPLICATION
machine_rate	Shop hourly burden rate, driven by region, equipment, and facility overhead	Supplier selection and region strategy are the primary lever. Rate variation is structural, not negotiable line-item.
machining_hours	Cycle time, driven by complexity, tolerances, setups, and process capability	Design for manufacturability and supplier process efficiency. This is the engineering-sourcing interface.
material_price	Raw material cost per pound, volatile for Ti and Inconel and stable for aluminum and steel	Material escalation clauses, long-term agreements, and strategic inventory carry real value on high-BTF parts.

scrap_factor

Process yield loss, driven by material machinability and supplier capability

Lower impact in most cases, but high-scrap materials (Ti, Inconel) warrant explicit yield negotiation.

07 — CONFIDENCE SCORING

The estimate tells you the number. The confidence score tells you how much to trust it.

Every estimate carries a confidence rating (High, Medium, or Low) assigned by a demerit system that evaluates input quality and inherent model uncertainty. The score is deterministic and fully traceable. It is not a gut check; it is a structured signal about where the estimate sits on the spectrum from "parametric baseline" to "specific enough to negotiate from."

HIGH — 0 DEMERITS

- All inputs from standard tables
- No overrides applied
- Standard or low-volatility material
- Complexity tier 1–4
- Use as: opening negotiation anchor

MEDIUM — 1-2 DEMERITS

- High-volatility material (Ti, Inconel)
- Unverified override supplied
- 3+ outside processes
- Use as: directional reference only
- Validated estimate before award

LOW — 3+ DEMERITS

- Multiple compounding factors
- Complexity tier 5
- Volatile material + overrides
- Use as: rough order of magnitude
- Supplier quotes required before action

Override semantics are explicit: an override tagged as **Verified** (sourced from a routing, CAM output, or supplier process sheet) carries no confidence penalty; it improves specificity. An override tagged as **Unverified** (a user estimate) reduces confidence by one demerit per field. The distinction prevents the confidence system from penalizing good data.

08 — OUTPUT EXAMPLE

Five parts across the material and complexity spectrum

The following outputs were generated from the live engine using the synthetic dataset. All part numbers, programs, and values are illustrative. Region adjusts machine rates, overhead rates, and margin assumptions. Geometry and hours are held constant across regions.

PART ID	DESCRIPTION	MATERIAL	COMPLEXITY	REGION	LOW	MID	HIGH	CONF.
P-3001	Aluminum Mounting Bracket	AL-6061	Tier 2	US	\$152	\$211	\$299	High
P-3002	Stainless Steel Valve Body	SS-17-4	Tier 3	US	\$614	\$886	\$1,299	High
P-3003	Titanium Actuator Housing	Ti-6Al-4V	Tier 4	US	\$1,572	\$2,206	\$3,158	Medium
P-3004	Inconel Combustion Liner	Inconel-718	Tier 5	US	\$3,072	\$4,295	\$6,120	Medium
P-3005	Steel Structural Fitting	Steel-4340	Tier 3	Asia	\$425	\$609	\$896	High

P-3003

\$2,206

\$1,572 - \$3,158

Titanium Actuator Housing (synthetic)

Ti-6Al-4V | Complexity 4 | Tight Tolerances | US | Vol: 1,500/yr | Batch: 250

PRICE BAND



COST BREAKDOWN - MID SCENARIO

<p>MATERIAL</p> <p>\$184.80</p> <p>8.4% of price</p>	<p>MACHINING</p> <p>\$1,085.00</p> <p>49.2% of price</p>	<p>OUTSIDE PROCESSES</p> <p>\$120.00</p> <p>5.4% of price</p>	<p>OVERHEAD & MARGIN (BURDENED ADDERS)</p> <p>\$772.43</p> <p>35.0% of price</p>
---	---	--	---

TOP COST DRIVERS - SENSITIVITY (+10% EACH)

MACHINE RATE	+\$172.23 (+7.8%)	supplier selection and region strategy
MACHINING HRS	+\$171.99 (+7.8%)	cycle time and process capability
MATERIAL PRICE	+\$29.29 (+1.3%)	Ti-6Al-4V market pricing

● Confidence: Medium. Ti-6Al-4V is a high-volatility material; material price range is wide.

All parts, part numbers, programs, and values are synthetic and illustrative only.

Key observation: machining cost accounts for 49% of unit price on a Ti-6Al-4V Tier 4 part. Machine rate and machining hours are nearly equal sensitivity drivers (+\$172 vs +\$172 on a +10% move). Material price, despite titanium's cost, contributes only +\$29. The negotiation lever is cycle time and supplier efficiency, not material markup. This is the insight a should-cost model surfaces before the first quote is received.

09 - FRAMEWORK CONTROLS

Four design decisions that make the model defensible under scrutiny

OVERHEAD BASIS TOGGLE

Overhead can be applied to all direct cost (default, treating overhead as a full SG&A and facility proxy) or to labor and machine cost only (treats material and outside processes as passthroughs). The active basis is recorded in every Excel output. The distinction matters when challenging a supplier's overhead allocation in a negotiation.

VERIFIED OVERRIDE SOURCE

Overrides to machining hours or setup hours can be tagged as Verified (from a routing, CAM output, or supplier process sheet) or Unverified (user estimate). Verified overrides improve estimate specificity with no confidence penalty. Unverified overrides reduce confidence by one demerit each. The system doesn't penalize good data.

PRICE FLOOR ENFORCEMENT

AI NARRATIVE LAYER

The low band cannot fall below a defensible minimum by complexity tier (Tier 1: \$25, Tier 3: \$200, Tier 5: \$1,500). This prevents optimistic assumptions from stacking into an implausible result. When the floor binds, the sensitivity output notes it: a muted sensitivity delta is a signal, not a clean result.

The top estimate can be passed through a structured Claude prompt that generates a plain-language sourcing summary covering cost drivers, band interpretation, and negotiation implication. The AI receives only computed outputs and generates explanation text only. It never produces, modifies, or overrides any cost number. Scoring is fully deterministic.

10 – LIMITATIONS & SCOPE

What this model does not do, and why that is an explicit decision

A should-cost model that doesn't acknowledge its own boundaries is more dangerous than no model at all. The following limitations are deliberate scope decisions, not oversights. Each has a corresponding V2 path if the use case demands it.

KNOWN

Machined parts only. Castings, forgings, sheet metal, composites, and additive are out of scope. The parametric assumptions (buy-to-fly, machine type, setup hours) are built for subtractive machining and do not translate to other manufacturing families.

KNOWN

Hours do not scale with part size. A 0.5 lb and a 15 lb part at the same complexity tier receive the same machining hour estimate. Size scaling via finished weight bands remains the highest-priority next improvement. Until then, the model is most reliable in the 0.5–10 lb range.

KNOWN

Outside process costs are per-piece benchmarks. Lot charges, minimum order quantities, and batch economics are not modeled. For high outside-process-cost parts at low volumes, actual quotes may exceed the model's per-piece estimate.

KNOWN

No tooling, NRE, or first-article costs. Non-recurring engineering, fixture fabrication, and first-article inspection costs are not included. The model estimates recurring unit cost only.

KNOWN

Python is the source of truth. The Excel export is a read-only output. Values edited in Excel do not feed back into the engine. This is an explicit V1 scope decision; round-trip editing would create stakeholder expectation drift around model ownership.

V2 ✓

STEP file geometry ingestion. Built and live. The geometry reader extracts bounding box, volume, surface area, and derived metrics (buy-to-fly ratio, slenderness ratio) directly from uploaded STEP files. These auto-populate engine inputs and can be overridden manually. Partially closes the size-scaling gap and moves the model toward geometry-aware estimation.

V2 ✓

Streamlit UI wrapper. Built and live at kenyon-should-cost.streamlit.app. The browser-based interface allows users to upload a STEP file, adjust inputs, run the estimate, view the results dashboard, and export to Excel without touching Python.

11 – BROADER CONTEXT

Third tool in a sourcing operations suite

The Should-Cost Model is the technical centerpiece of a five-tool sourcing operations suite. Each tool is standalone, solves a specific operational problem, and produces outputs defensible in a leadership setting. The suite is designed for the sourcing professional who needs to present structured decisions to leadership, not for the analyst building dashboards.

SUITE PROGRESS

DESIGN PHILOSOPHY

- Part Prioritization Framework – Complete
- RFQ Lifecycle Tracker – Complete
- Parametric Should-Cost Model – Complete
- Supplier Performance & Risk Scorecard – Planned
- Make vs. Buy Decision Framework – Planned

Clarity over complexity. Every number traces back to an assumption. Every assumption is editable. The model surfaces the estimate. The sourcing engineer makes the call. Built in Python with Excel export. No enterprise tooling required, no black boxes, no licensing dependency.

12 – LIVE TOOL

Try the model live

`kenyon-should-cost.streamlit.app`

Enter a part, select material and complexity, upload a STEP file or enter part inputs manually. The engine generates a three-scenario price band, sensitivity analysis, and AI narrative in seconds. Excel export included.

[OPEN TOOL →](#)

13 – CLOSING NOTE

Built to close the gap between "I have no idea what this should cost" and "I have a defensible number to negotiate from." Parametric by design, explainable by requirement, and modular enough to extend when the use case demands it. **The model produces the estimate. The sourcing engineer decides what to do with it.**

Python cost engine + Excel export + Streamlit UI. | 8 materials, 5 complexity tiers, 3 regions, 9 outside processes. | Three-scenario price band (low / mid / high). | Sensitivity analysis, confidence scoring, AI narrative layer. | STEP geometry ingestion with buy-to-fly derivation. | Overhead basis toggle, verified override semantics, price floor enforcement. | 26 passing tests. | Synthetic dataset benchmarked against US aerospace machining benchmarks.