

The Capers-Workman Equation

A worksheet for comparing the cost of AI alone with the cost of AI-powered people, by Valerie Capers Workman, Esq.

SAMPLE WORKSHEET · PROPOSED MODEL FOR PRACTITIONER TESTING

The sample workflow: an AI tool reviews incoming commercial contracts (vendor and customer agreements), flags terms that deviate from the company's standard positions, and drafts a redline for attorney sign-off.

Illustrative example. The figures are invented to show how the worksheet works. They are not data from any organization.

Cost per unit of useful work = total cost ÷ units of useful work delivered

Run it before you choose a path. Use projected costs and output for one workflow, over the same time horizon, at the same quality bar. The lower cost per unit wins.

$$C_A = \frac{L + I + H + Q}{U_A}$$

$$C_H = \frac{S + D + A - R - K}{U_H}$$

FIRST, DEFINE THE UNIT AND THE PERIOD

Workflow

AI review of commercial contracts

Unit of useful work

Contracts reviewed, redlined, ready to sign

Decision horizon (e.g. first 12 months after launch)

First 12 months after launch

Option A: AI alone

The AI reviews each contract and drafts the redline. Attorneys verify it and sign off.

TERM		PROJECTED VALUE	BASIS
L	License or amortized build License quote, or build cost spread across the horizon	\$ 180,000	Quote · Bench Data · Assume
I	Inference at production volume Projected full-scale usage at the expected rate, not pilot pricing	\$ 72,000	Quote · Bench Data · Assume
H	Human verification People needed to review, approve, or correct AI output before it is used	\$ 260,000	Quote · Bench Data · Assume
Q	Quality and remediation Expected cost of errors that get through: rework, escalations, customer and compliance costs	\$ 95,000	Quote · Bench Data · Assume
Total cost		607,000	
U_A	Units of useful work AI delivers Projected units the AI will get right and that hold	9,000	Quote · Bench Data · Assume

Cost per unit = \$67.44

Option B: AI-powered people

Contract reviewers do the review, using the same AI tool to work faster.

TERM		PROJECTED VALUE	BASIS
S	Loaded salary and benefits Salary plus benefits and employer taxes for the people in the workflow	\$ 620,000	Quote · Bench Data · Assume
D	Employee development Planned training and ramp time to work well with AI	\$ 40,000	Quote · Bench Data · Assume
A	AI tools and inference Seats or amortized build, plus the inference people use	\$ 24,000	Quote · Bench Data · Assume
R	Attrition avoided SUBTRACT Expected turnover cost this path avoids: recruiting, onboarding, lost ramp time	\$ 45,000	Quote · Bench Data · Assume
K	Knowledge retained SUBTRACT Value of expertise kept by retaining experienced people	\$ 0	Quote · Bench Data · Assume

Count each benefit once. If retained people or knowledge already show up as higher projected output in U_H, enter 0 for R or K.

Total cost	639,000	
U_H	Units of useful work people deliver Current output plus the expected lift from AI, at the same quality bar as Option A	11,500

Cost per unit = \$55.57

The result

AI alone: $\$607,000 \div 9,000 = \$67.44$ per contract

AI-powered people: $\$639,000 \div 11,500 = \text{\textbf{\$55.57 per contract}}$

AI-powered people cost 18% less per unit. The AI path looks cheaper until human verification (H) and remediation (Q) are counted. Together they add \$355,000.

The stress test

Least certain estimate on the winning side: the productivity lift behind U_H , assumed at 20%.

Worst case, 5% lift: $\$639,000 \div 10,080 = \$63.39$ per contract

Least certain estimate on the other side: Q, an assumption. At a best case of \$40,000, AI alone falls to \$61.33.

The winner holds in both tests. Before committing, firm up the lift estimate with a short trial.

Assumptions behind the sample

- Four reviewers at \$155,000 loaded cost each deliver 9,600 contracts a year today.
- With AI tools, their output rises an expected 20%, to about 11,500.
- The AI drafts 12,000 reviews a year. Benchmarks suggest 75% meet the standard and hold, so U_A is 9,000.
- Every AI review needs attorney sign-off: two reviewers at \$130,000 loaded cost each (H).
- K is 0 because retained expertise already shows up in the projected output. Counting it again would count it twice.

Where the estimates come from

You do not need to implement anything to run the equation. Every value can be projected from sources you already have. Circle the basis for each value in the Basis column:

- **Quote:** vendor pricing, build estimates, and training costs.
- **Bench (benchmark):** published results, industry data, or comparable deployments.
- **Data (your own):** current salaries, output, error rates, and turnover for this workflow.
- **Assume (assumption):** a stated judgment call. Write it down so others can challenge it.

Steps for Option A: AI alone

1. Price the AI (L): the license quote, or the build cost amortized across the horizon.
2. Project inference (I) at full production volume, at the expected rate, not pilot pricing.
3. Add the costs usually left out: human verification (H) and quality and remediation (Q).
4. Estimate output (U_A): the units the AI will get right and that hold, based on benchmarks, comparable deployments, or a short trial.

Steps for Option B: AI-powered people

1. Enter the loaded salary and benefits (S) from your current data.
2. Add planned development (D) and the AI people will use (A): seats or amortized build, plus inference.
3. Subtract expected attrition avoided (R) and knowledge retained (K), only where they do not already show up as higher output.
4. Estimate output (U_H): current output for this workflow plus the expected lift from AI, at the same quality bar.

Then compare and stress test

Divide each option's total cost by its output. The lower cost per unit wins. Then rerun the equation with your least certain estimate at its worst case. If the winner does not change, decide with confidence. If it does, that one input decides the outcome. Firm it up with a quote, a benchmark, or a short trial of a few weeks before you commit. After you choose a path, track actual results against the estimates to improve your next analysis.

The equation does not assume people or AI will win. It measures which operating model delivers useful work at the lower total cost, and under what conditions.

The measurement of attrition avoided and knowledge retained remains open to validation.

© 2026 Valerie Capers Workman, Esq. · The Capers-Workman Equation™