Unit Economics Worksheet

The Unit Economics Worksheet converts the cost inventory in e6-cost-model into a per-offer gross-margin model by comparing fully-loaded operating cost to the user value or revenue each offer generates. This worksheet is the C12 cost source: c10-revenue-model-design must point every pricing hypothesis to a populated row here via unit_cost_ref, and the anti-vaporware rule (C12) declares that a pricing hypothesis is invalid unless the referenced row is populated and its margin meets the min_margin_pct floor. Rows that show automation increasing cost rather than reducing it must be flagged explicitly — hiding negative-margin automation is a common failure mode in early agent deployments.

Worksheet Structure

Each offer row captures the fully-loaded unit cost, the revenue or value realization per unit, the gross margin, and the break-even volume. The worksheet is populated from e6-cost-model cost-per-outcome figures; do not invent figures not traceable to that spec.

Per-offer row schema

```
- offer_id: <kebab-case; used as unit_cost_ref in c10-revenue-model-design>
  offer_name: <short descriptive name>
  offer_type: <task | workflow | user-seat | outcome>

  # Cost side (from e6-cost-model)
  cost_inputs:
    model_api_cost_per_unit_usd: <number>
    compute_cost_per_unit_usd: <number>
    storage_retrieval_cost_per_unit_usd: <number>
    tooling_cost_per_unit_usd: <number>
    human_review_cost_per_unit_usd: <number — HITL time-value; must not be zero if HITL
gates exist>
    egress_observability_cost_per_unit_usd: <number>
    total_cogs_per_unit_usd: <derived — sum of above>

  # Revenue / value side
  revenue_inputs:
    revenue_per_unit_usd: <number — price charged or value attributed per unit>
    revenue_model: <subscription-allocated | usage-fee | outcome-fee | internal-cost-
avoidance>
    units_per_month: <number — current or projected volume>

  # Margin
  gross_margin_pct: <derived — (revenue_per_unit - total_cogs_per_unit) / revenue_per_unit
× 100>
  gross_margin_usd_per_unit: <derived>
  monthly_gross_profit_usd: <derived — gross_margin_usd_per_unit × units_per_month>

  # Break-even
  fixed_cost_allocation_usd: <monthly fixed costs allocated to this offer>
  break_even_units_per_month: <derived — fixed_cost / gross_margin_usd_per_unit>

  # Flags
  automation_increases_cost: <true | false — set true if total_cogs_per_unit >
manual_baseline_cost_per_unit>
  manual_baseline_cost_per_unit_usd: <number — cost of the same task done manually;
required if automation_increases_cost=true>
  viable: <true | false — false if gross_margin_pct < min_margin_pct or
automation_increases_cost=true without a strategic justification>
  viability_note: <required if viable=false>
```

```
  min_margin_pct: <number — the floor below which the offer is not viable per C12>
  last_updated: <ISO-8601>
```

Gross-Margin Interpretation Guide

Gross margin range	Interpretation	Action
> 70%	Strong unit economics — invest in volume growth	Confirm cost model is fully-loaded (human review included)
50–70%	Healthy for software-delivered services	Monitor HITL frequency; right-size infra
30–50%	Acceptable for complex AI-augmented workflows	Identify top COGS driver; target waste reduction
10–30%	Thin — vulnerable to model price changes or volume spikes	Address before scaling; do not add fixed costs
< 10%	Marginal — not viable as a standalone product	Flag viable: false; strategic justification required
Negative	Automation increases cost vs. manual baseline	Flag automation_increases_cost: true; do not scale

Worked Example: GROW Library Stormwater Pre-Screener (illustrative)

The GoodSam.ai stormwater pre-screener helps applicants in South Walton County self-assess whether a proposed project triggers stormwater review under FDEP delegated authority. All figures are (illustrative). The offer is priced as a per-screening fee to special districts.

```
- offer_id: stormwater-pre-screener-screening
  offer_name: Stormwater Pre-Screener — per screening
  offer_type: task

  cost_inputs:
    model_api_cost_per_unit_usd: 0.047      # (illustrative) ~2 model calls: cheap classify
+ standard narrative
    compute_cost_per_unit_usd: 0.003        # (illustrative) cloud function
    storage_retrieval_cost_per_unit_usd: 0.018  # (illustrative) embeddings + vector search
    tooling_cost_per_unit_usd: 0.009        # (illustrative) GIS parcel lookup + public
records check
    human_review_cost_per_unit_usd: 0.052   # (illustrative) 8 min × $65/hr × 6% escalation
rate
    egress_observability_cost_per_unit_usd: 0.002   # (illustrative)
    total_cogs_per_unit_usd: 0.131          # (illustrative)

  revenue_inputs:
    revenue_per_unit_usd: 0.75              # (illustrative) per-screening fee to district
    revenue_model: usage-fee
    units_per_month: 620                    # (illustrative) South Walton + 2 neighboring
districts

  gross_margin_pct: 82.5                   # (illustrative)
  gross_margin_usd_per_unit: 0.619         # (illustrative)
  monthly_gross_profit_usd: 383.8          # (illustrative)
```

```
  fixed_cost_allocation_usd: 95            # (illustrative) shared infra + support
allocation
  break_even_units_per_month: 154          # (illustrative)

  automation_increases_cost: false
  manual_baseline_cost_per_unit_usd: 2.85  # (illustrative) ~26 min clerk time per manual
triage
  viable: true
  viability_note: ""
  min_margin_pct: 50
  last_updated: 2026-05-29
```

Second offer: clerk-queue routing workflow (illustrative)

```
- offer_id: stormwater-pre-screener-workflow
  offer_name: Stormwater Pre-Screener — monthly district seat
  offer_type: user-seat

  cost_inputs:
    model_api_cost_per_unit_usd: 12.40     # (illustrative) all screenings bundled for a
seat
    compute_cost_per_unit_usd: 1.20
    storage_retrieval_cost_per_unit_usd: 4.80
    tooling_cost_per_unit_usd: 2.90
    human_review_cost_per_unit_usd: 18.50  # (illustrative) HITL is dominant at seat level
    egress_observability_cost_per_unit_usd: 0.80
    total_cogs_per_unit_usd: 40.60         # (illustrative)

  revenue_inputs:
    revenue_per_unit_usd: 199.00           # (illustrative) monthly seat license to district
    revenue_model: subscription-allocated
    units_per_month: 4                     # (illustrative) four district seats

  gross_margin_pct: 79.6                  # (illustrative)
  gross_margin_usd_per_unit: 158.4        # (illustrative)
  monthly_gross_profit_usd: 633.6         # (illustrative)

  fixed_cost_allocation_usd: 180          # (illustrative)
  break_even_units_per_month: 2           # (illustrative)

  automation_increases_cost: false
  manual_baseline_cost_per_unit_usd: 210.00  # (illustrative) ~3.2 hr clerk time per month
per district seat
  viable: true
  viability_note: ""
  min_margin_pct: 50
  last_updated: 2026-05-29
```

Example of a negative-margin offer: automated public-comment drafting (illustrative)

```
- offer_id: public-comment-auto-draft
  offer_name: Automated Public Comment Draft — per draft
  offer_type: task

  cost_inputs:
    model_api_cost_per_unit_usd: 0.210     # (illustrative) premium tier required; cites
statutes
    compute_cost_per_unit_usd: 0.008
    storage_retrieval_cost_per_unit_usd: 0.035
    tooling_cost_per_unit_usd: 0.012
    human_review_cost_per_unit_usd: 0.325  # (illustrative) attorney review required;
expensive
    egress_observability_cost_per_unit_usd: 0.004
    total_cogs_per_unit_usd: 0.594         # (illustrative)
```

```
  revenue_inputs:
    revenue_per_unit_usd: 0.50            # (illustrative) proposed per-draft fee
    revenue_model: usage-fee
    units_per_month: 80

  gross_margin_pct: -18.8                 # (illustrative) NEGATIVE
  gross_margin_usd_per_unit: -0.094
  monthly_gross_profit_usd: -7.52

  fixed_cost_allocation_usd: 30
  break_even_units_per_month: 0           # unreachable at current price

  automation_increases_cost: true
  manual_baseline_cost_per_unit_usd: 0.45  # (illustrative) manual draft is cheaper at this
volume
  viable: false
  viability_note: "Attorney-review HITL gate makes per-draft cost exceed proposed price. Do
not scale at $0.50 per draft. Either raise price to $1.20+, reduce HITL frequency via
quality improvement, or move to a subscription model that amortizes fixed attorney-review
time. Flag for c10-revenue-model-design review before any pricing hypothesis references
this offer_id."
  min_margin_pct: 50
  last_updated: 2026-05-29
```

This third row demonstrates the required automation_increases_cost: true flag. The c10-revenue-model-design anti-vaporware rule (C12) explicitly blocks a pricing hypothesis that cites unit_cost_ref: public-comment-auto-draft while that row carries viable: false — the commercialization team cannot claim a margin that the cost model does not support.

C12 Linkage Notes

This worksheet is the e6-unit-economics-worksheet read by c10-revenue-model-design. The C12 contract requires:

1. Every unit_cost_ref in a pricing hypothesis resolves to a real offer_id row in this worksheet.

2. The resolved row must have viable: true.

3. The row's gross_margin_pct must meet or exceed the pricing hypothesis's min_margin_pct.

4. Any row with automation_increases_cost: true must carry a viability_note before it can be cited in a pricing hypothesis.

A pricing hypothesis that fails any of these checks is invalid and must be revised before commercialization work proceeds. This worksheet is updated by the e6-cost-model owner on each billing period; c10-revenue-model-design readers must check last_updated before relying on margin figures.