

The AI pricing playbook for founders

The fundamental shift → AI economics ≠ SaaS economics. Every AI query costs money (compute, inference, human-in-the-loop). You'll see 50-60% gross margins vs. 80-90% for traditional SaaS. Microsoft lost \$20/user/month on GitHub Copilot initially. If your unit economics don't work at 10 customers, they won't magically work at 1,000.

Three business models, three pricing approaches

Model	What it does	How to price it
Copilots	Assist humans, boost productivity	Per seat or consumption (like SaaS)
Agents	Execute entire workflows autonomously	Outcome/workflow-based (tied to tangible ROI)
AI-enabled Services	Blend automation + human oversight	Consumption to outcome-based; compare to FTE cost

Your charge metric = your strategy

Three options, one critical trade-off: As you move from consumption → workflow → outcome-based pricing, you accept more cost risk for tighter value alignment.

1. Consumption (tokens/API calls): Clean margins, predictable costs. Problem: Customers don't think in tokens. Works only for technical buyers.
2. Workflow (per task completed): Closer to how work happens. Clearer value prop, higher cost variability. Works when tasks are the outcome.
3. Outcome (per result delivered): Maximum value alignment, maximum cost risk. Intercom's \$0.99/resolved ticket is the gold standard. Works when you're confident in AI performance and outcomes are measurable.

Winning move: Hybrid models (base subscription + usage/outcome tiers) provide predictability while capturing upside. Examples: EvenUp, Leena AI, Sett.ai, Zenskar.

The ROI reality check

Map yourself honestly on two dimensions:

Value type: Revenue uplift vs. cost savings

ROI clarity: Hard (measurable) vs. soft (incremental)

- Copilots = Soft ROI = Weak pricing power. Customers question "are we getting value?"
- Agents = Hard ROI = Strong pricing power. They close the loop, outcomes are undeniable.
- Service replacements = Clear cost reduction. Master the TCO comparison—enterprises underestimate legacy costs.

2026 warning: Soft ROI worked in 2025's "AI adoption at all costs" environment. As pilots hit renewal, pricing must reflect actual value delivered, not promise.

Five founder imperatives

1. Lead with value, not costs

Don't do cost-plus (calculate costs, double it). Frame it as: "We're creating \$X in value, charging Y%." Most founders undercharge because asking for more feels awkward.

Hybrid formula that works:

- Platform fee (2X delivery costs) + outcome credits
- Example: \$12K/year covers infrastructure + 100 ticket resolutions; \$5K per additional 100 tickets

2. Find your sweet spot through friction

- Start with a price (\$12K/year)
- If customers say "sold" immediately → too cheap
- Raise incrementally until you hear "we need to think about that"
- Stop just before it becomes a blocker

This is how multi-billion dollar companies found their sweet spot.

3. Build unit economics discipline from day one

Track true costs:

- CEO time on sales = sales cost
- CTO time on support = support cost
- Full cost stack: inference + human-in-the-loop + CS overhead

Companies that ignore this scale to negative margins without realizing it.

4. Avoid the pricing complexity trap

Nine different pricing approaches across contracts becomes an operational nightmare at scale. Identify one model that works at both 10 and 1,000 customers.

5. Align pricing with your entire GTM

Pricing determines how sales, CS, and product teams operate. Outcome-based contracts create new questions:

- How do AEs size scalable deals?
- How much expansion upside should CS capture?

Your pricing model becomes your organizational North Star.

Critical questions before launch

1. Can customers understand what they're paying for without a manual?
2. Does your pricing reflect compute costs AND customer value?
3. Are you measuring the right success metrics? (AI resolution rate, % work autonomous, developer acceptance rate—not just ARR/CAC)
4. Can you absorb cost variability if you go outcome-based?
5. Does pricing work for 10 customers AND 1,000 customers?

The bottom line

AI doesn't monetize access. It monetizes outcomes. Companies like Intercom, EvenUp, and Leena AI align every team around delivered work: tickets resolved, documents completed, problems solved. The winners will charge for what their AI earns, not what it costs or what customers access. Your charge metric isn't a billing decision—it's a statement about what you believe your AI is worth and what you're willing to stake your margins on to prove it. **The playbook is still being written. Move fast, test value early, price boldly.**